



MIPS® Architecture For Programmers

Volume I-B: Introduction to the microMIPS64® Architecture

Document Number: MD00743

Revision 5.03

Sept. 9, 2013

**MIPS Technologies, Inc.
955 East Arques Avenue
Sunnyvale, CA 94085-4521**

Copyright © 2001-2003,2005,2008-2013 MIPS Technologies Inc. All rights reserved.

Unpublished rights (if any) reserved under the copyright laws of the United States of America and other countries.

This document contains information that is proprietary to MIPS Technologies, Inc. ("MIPS Technologies") one of the Imagination Technologies Group plc companies. Any copying, reproducing, modifying or use of this information (in whole or in part) that is not expressly permitted in writing by MIPS Technologies or an authorized third party is strictly prohibited. At a minimum, this information is protected under unfair competition and copyright laws. Violations thereof may result in criminal penalties and fines.

Any document provided in source format (i.e., in a modifiable form such as in FrameMaker or Microsoft Word format) is subject to use and distribution restrictions that are independent of and supplemental to any and all confidentiality restrictions. UNDER NO CIRCUMSTANCES MAY A DOCUMENT PROVIDED IN SOURCE FORMAT BE DISTRIBUTED TO A THIRD PARTY IN SOURCE FORMAT WITHOUT THE EXPRESS WRITTEN PERMISSION OF MIPS TECHNOLOGIES, INC.

MIPS Technologies reserves the right to change the information contained in this document to improve function, design or otherwise. MIPS Technologies does not assume any liability arising out of the application or use of this information, or of any error or omission in such information. Any warranties, whether express, statutory, implied or otherwise, including but not limited to the implied warranties of merchantability or fitness for a particular purpose, are excluded. Except as expressly provided in any written license agreement from MIPS Technologies or an authorized third party, the furnishing of this document does not give recipient any license to any intellectual property rights, including any patent rights, that cover the information in this document.

The information contained in this document shall not be exported, re-exported, transferred, or released, directly or indirectly, in violation of the law of any country or international law, regulation, treaty, Executive Order, statute, amendments or supplements thereto. Should a conflict arise regarding the export, re-export, transfer, or release of the information contained in this document, the laws of the United States of America shall be the governing law.

The information contained in this document constitutes one or more of the following: commercial computer software, commercial computer software documentation, or other commercial items. If the user of this information, or any related documentation of any kind, including related technical data or manuals, is an agency, department, or other entity of the United States government ("Government"), the use, duplication, reproduction, release, modification, disclosure, or transfer of this information, or any related documentation of any kind, is restricted in accordance with Federal Acquisition Regulation 12.212 for civilian agencies and Defense Federal Acquisition Regulation Supplement 227.7202 for military agencies. The use of this information by the Government is further restricted in accordance with the terms of the license agreement(s) and/or applicable contract terms and conditions covering this information from MIPS Technologies or an authorized third party.

MIPS, MIPS I, MIPS II, MIPS III, MIPS IV, MIPS V, MIPSr3, MIPS32, MIPS64, microMIPS32, microMIPS64, MIPS-3D, MIPS16, MIPS16e, MIPS-Based, MIPSsim, MIPSpro, MIPS-VERIFIED, Aptiv logo, microAptiv logo, interAptiv logo, microMIPS logo, MIPS Technologies logo, MIPS-VERIFIED logo, proAptiv logo, 4K, 4Kc, 4Km, 4Kp, 4KE, 4KEc, 4KEm, 4KEp, 4KS, 4KSc, 4KSd, M4K, M14K, 5K, 5Kc, 5Kf, 24K, 24Kc, 24Kf, 24KE, 24KEc, 24KEf, 34K, 34Kc, 34Kf, 74K, 74Kc, 74Kf, 1004K, 1004Kc, 1004Kf, 1074K, 1074Kc, 1074Kf, R3000, R4000, R5000, Aptiv, ASMACRO, Atlas, "At the core of the user experience.", BusBridge, Bus Navigator, CLAM, CorExtend, CoreFPGA, CoreLV, EC, FPGA View, FS2, FS2 FIRST SILICON SOLUTIONS logo, FS2 NAVIGATOR, HyperDebug, HyperJTAG, IASim, iFlowtrace, interAptiv, JALGO, Logic Navigator, Malta, MDMX, MED, MGB, microAptiv, microMIPS, Navigator, OCI, PDtrace, the Pipeline, proAptiv, Pro Series, SEAD-3, SmartMIPS, SOC-it, and YAMON are trademarks or registered trademarks of MIPS Technologies, Inc. in the United States and other countries.

All other trademarks referred to herein are the property of their respective owners.

Template: nB1.03, Built with tags: 2B ARCH FPU_PS FPU_PSandARCH MIPS64

Contents

Chapter 1: About This Book	10
1.1: Typographical Conventions	10
1.1.1: Italic Text	11
1.1.2: Bold Text	11
1.1.3: Courier Text	11
1.2: UNPREDICTABLE and UNDEFINED	11
1.2.1: UNPREDICTABLE	11
1.2.2: UNDEFINED	12
1.2.3: UNSTABLE	12
1.3: Special Symbols in Pseudocode Notation	12
1.4: For More Information	15
 Chapter 2: The MIPS Architecture: An Introduction	 16
2.1: MIPS Instruction Set Overview	16
2.1.1: Historical Perspective	16
2.1.2: Architectural Evolution	17
2.1.3: Architectural Changes Relative to the MIPS I through MIPS V Architectures	21
2.2: Compliance and Subsetting	21
2.3: Components of the MIPS Architecture	24
2.3.1: MIPS Instruction Set Architecture (ISA)	24
2.3.2: MIPS Privileged Resource Architecture (PRA)	24
2.3.3: MIPS Modules and Application Specific Extensions (ASEs)	24
2.3.4: MIPS User Defined Instructions (UDIs)	24
2.4: Architecture Versus Implementation	24
2.5: Relationship between the MIPSr3 Architectures	25
2.6: Pipeline Architecture	27
2.6.1: Pipeline Stages and Execution Rates	27
2.6.2: Parallel Pipeline	28
2.6.3: Superpipeline	28
2.6.4: Superscalar Pipeline	29
2.7: Load/Store Architecture	29
2.8: Programming Model	30
2.8.1: CPU Data Formats	30
2.8.2: FPU Data Formats	30
2.8.3: Coprocessors (CP0-CP3)	31
2.8.4: CPU Registers	31
2.8.5: FPU Registers	33
2.8.6: Byte Ordering and Endianness	37
2.8.7: Memory Alignment	39
2.8.8: Memory Access Types	40
2.8.9: Implementation-Specific Access Types	42
2.8.10: Cacheability and Coherency Attributes and Access Types	42
2.8.11: Mixing Access Types	43
2.8.12: Instruction Fetches	43
 Chapter 3: Modules and Application Specific Extensions	 48
3.1: Description of Optional Components	48

3.2: List of Application Specific Instructions	49
3.2.1: The MDMX™ Application Specific Extension to the microMIPS64 Architectures	49
3.2.2: The MIPS-3D® Application Specific Extension to the microMIPS Architecture	50
3.2.3: The SmartMIPS® Application Specific Extension to the microMIPS32 Architecture	50
3.2.4: The MIPS® DSP Module to the MIPS Architecture	50
3.2.5: The MIPS® MT Module to the MIPS Architecture.....	50
3.2.6: The MIPS® MCU Application Specific Extension to the MIPS Architecture	50
3.2.7: The MIPS® Virtualization Module to the MIPS Architecture	50
3.2.8: The MIPS® SIMD Architecture Module to the MIPS Architecture.....	50
Chapter 4: Overview of the CPU Instruction Set	52
4.1: CPU Instructions, Grouped By Function.....	52
4.1.1: CPU Load and Store Instructions.....	52
4.1.2: Computational Instructions.....	56
4.1.3: Jump and Branch Instructions.....	61
4.1.4: Miscellaneous Instructions	63
4.1.5: Coprocessor Instructions	66
4.1.6: CPU Instruction Restrictions	67
Chapter 5: Overview of the FPU Instruction Set.....	68
5.1: Binary Compatibility	68
5.2: Enabling the Floating Point Coprocessor	69
5.3: IEEE Standard 754.....	69
5.4: FPU Data Types	69
5.4.1: Floating Point Formats	69
5.4.2: Fixed Point Formats	74
5.5: Floating Point Register Types	74
5.5.1: FPU Register Models	75
5.5.2: Binary Data Transfers (32-Bit and 64-Bit)	75
5.5.3: FPRs and Formatted Operand Layout.....	76
5.6: Floating Point Control Registers (FCRs)	77
5.6.1: Floating Point Implementation Register (FIR, CP1 Control Register 0)	77
5.6.2: User Floating point Register mode control (UFR, CP1 Control Register 1).....	80
5.6.3: User Negated Floating point Register mode control (UNFR, CP1 Control Register 4).....	80
5.6.4: Floating Point Control and Status Register (FCSR, CP1 Control Register 31).....	81
5.6.5: Floating Point Condition Codes Register (FCCR, CP1 Control Register 25).....	85
5.6.6: Floating Point Exceptions Register (FEXR, CP1 Control Register 26)	85
5.6.7: Floating Point Enables Register (FENR, CP1 Control Register 28).....	86
5.7: Formats of Values Used in FP Registers	87
5.8: FPU Exceptions.....	87
5.8.1: Exception Conditions	88
5.9: FPU Instructions	92
5.9.1: Data Transfer Instructions.....	92
5.9.2: Arithmetic Instructions.....	94
5.9.3: Conversion Instructions.....	95
5.9.4: Formatted Operand-Value Move Instructions	96
5.9.5: Conditional Branch Instructions	97
5.9.6: Miscellaneous Instructions	98
5.10: Valid Operands for FPU Instructions	98
5.11: FPU Instruction Formats.....	100
Appendix A: Revision History	102

Figures

Figure 2.1: MIPS Architectures	17
Figure 2.2: Relationship of the Binary Representations of MIPSr3 Architectures	26
Figure 2.3: Relationships of the Assembler Source Code Representations of the MIPSr3 Architectures	27
Figure 2.4: One-Deep Single-Completion Instruction Pipeline.....	28
Figure 2.5: Four-Deep Single-Completion Pipeline.....	28
Figure 2.6: Four-Deep Superpipeline.....	29
Figure 2.7: Four-Way Superscalar Pipeline	29
Figure 2.8: CPU Registers	33
Figure 2.9: FPU Registers for a 32-bit FPU	35
Figure 2.10: FPU Registers for a 64-bit FPU if Status _{FR} is 1	36
Figure 2.11: FPU Registers for a 64-bit FPU if Status _{FR} is 0.....	37
Figure 2.12: Big-Endian Byte Ordering	38
Figure 2.13: Little-Endian Byte Ordering.....	38
Figure 2.14: Big-Endian Data in Doubleword Format	39
Figure 2.15: Little-Endian Data in Doubleword Format.....	39
Figure 2.16: Big-Endian Misaligned Word Addressing.....	40
Figure 2.17: Little-Endian Misaligned Word Addressing	40
Figure 2.18: Three instructions placed in a 64-bit wide, little-endian memory	44
Figure 2.19: Three instructions placed in a 64-bit wide, big-endian memory	44
Figure 3.1: microMIPS ISAs and ASEs.....	49
Figure 5.1: Single-Precisions Floating Point Format (S).....	71
Figure 5.2: Double-Precisions Floating Point Format (D)	71
Figure 5.3: Paired Single Floating Point Format (PS)	71
Figure 5.4: Word Fixed Point Format (W)	74
Figure 5.5: Longword Fixed Point Format (L)	74
Figure 5.6: FPU Word Load and Move-to Operations	76
Figure 5.7: FPU Doubleword Load and Move-to Operations	76
Figure 5.8: Single Floating Point or Word Fixed Point Operand in an FPR	77
Figure 5.9: Double Floating Point or Longword Fixed Point Operand in an FPR.....	77
Figure 5.10: Paired-Single Floating Point Operand in an FPR	77
Figure 5.11: FIR Register Format	78
Figure 5.12: UFR Register Format	80
Figure 5.13: UNFR Register Format	81
Figure 5.14: FCSR Register Format	82
Figure 5.15: FCCR Register Format	85
Figure 5.16: FEXR Register Format	86
Figure 5.17: FENR Register Format	87

Tables

Table 1.1: Symbols Used in Instruction Operation Statements.....	12
Table 2.1: Unaligned Load and Store Instructions	40
Table 2.2: Speculative instruction fetches	45
Table 4.1: Load and Store Operations Using Register + Offset Addressing Mode.....	53
Table 4.2: FPU Load and Store Operations Using Register + Register Addressing Mode.....	53
Table 4.3: Aligned CPU Load/Store Instructions.....	54
Table 4.4: Unaligned CPU Load and Store Instructions	55
Table 4.5: Atomic Update CPU Load and Store Instructions	55
Table 4.6: CPU Load and Store Instructions Using Register + Register Addressing	56
Table 4.7: Coprocessor Load and Store Instructions.....	56
Table 4.8: FPU Load and Store Instructions Using Register + Register Addressing.....	56
Table 4.9: ALU Instructions With a 16-bit Immediate Operand.....	57
Table 4.10: Other ALU Instructions With a Immediate Operand.....	58
Table 4.11: Three-Operand ALU Instructions	58
Table 4.12: Two-Operand ALU Instructions.....	59
Table 4.13: Shift Instructions	59
Table 4.14: Multiply/Divide Instructions.....	60
Table 4.15: Unconditional Jump Within a 256 Megabyte Region	62
Table 4.16: Unconditional Jump using Absolute Address.....	62
Table 4.17: PC-Relative Conditional Branch Instructions Comparing Two Registers.....	63
Table 4.18: PC-Relative Conditional Branch Instructions Comparing With Zero	63
Table 4.19: PC-relative Unconditional Branch	63
Table 4.20: Serialization Instruction	64
Table 4.21: System Call and Breakpoint Instructions	64
Table 4.22: Trap-on-Condition Instructions Comparing Two Registers	64
Table 4.24: CPU Conditional Move Instructions	65
Table 4.25: Prefetch Instructions	65
Table 4.23: Trap-on-Condition Instructions Comparing an Immediate Value	65
Table 4.26: NOP Instructions.....	66
Table 4.27: Coprocessor Definition and Use in the MIPS Architecture.....	66
Table 5.1: Parameters of Floating Point Data Types	70
Table 5.2: Value of Single or Double Floating Point DataType Encoding.....	71
Table 5.3: Value Supplied When a New Quiet NaN Is Created	73
Table 5.4: FIR Register Field Descriptions	78
Table 5.5: UFR Register Field Descriptions.....	80
Table 5.6: UNFR Register Field Descriptions	81
Table 5.7: FCSR Register Field Descriptions	82
Table 5.8: Cause, Enable, and Flag Bit Definitions.....	84
Table 5.10: FCCR Register Field Descriptions	85
Table 5.9: Rounding Mode Definitions	85
Table 5.11: FEXR Register Field Descriptions.....	86
Table 5.12: FENR Register Field Descriptions	87
Table 5.13: Default Result for IEEE Exceptions Not Trapped Precisely	89
Table 5.14: FPU Data Transfer Instructions.....	93
Table 5.15: FPU Loads and Stores Using Register+Offset Address Mode	93
Table 5.16: FPU Loads and Using Register+Register Address Mode.....	93
Table 5.17: FPU Move To and From Instructions	93

Table 5.18: FPU IEEE Arithmetic Operations	94
Table 5.19: FPU-Approximate Arithmetic Operations	95
Table 5.20: FPU Multiply-Accumulate Arithmetic Operations	95
Table 5.21: FPU Conversion Operations Using the <i>FCSR</i> Rounding Mode.....	95
Table 5.23: FPU Formatted Operand Move Instructions	96
Table 5.22: FPU Conversion Operations Using a Directed Rounding Mode	96
Table 5.24: FPU Conditional Move on True/False Instructions.....	97
Table 5.25: FPU Conditional Move on Zero/Nonzero Instructions.....	97
Table 5.26: FPU Conditional Branch Instructions	98
Table 5.27: CPU Conditional Move on FPU True/False Instructions	98
Table 5.28: FPU Operand Formats.....	98
Table 5.29: Valid Formats for FPU Operations	99

About This Book

The MIPS® Architecture For Programmers Volume I-B: Introduction to the microMIPS64® Architecture comes as part of a multi-volume set.

- Volume I-A describes conventions used throughout the document set, and provides an introduction to the MIPS64® Architecture
- Volume I-B describes conventions used throughout the document set, and provides an introduction to the microMIPS64™ Architecture
- Volume II-A provides detailed descriptions of each instruction in the MIPS64® instruction set
- Volume II-B provides detailed descriptions of each instruction in the microMIPS64™ instruction set
- Volume III describes the MIPS64® and microMIPS64™ Privileged Resource Architecture which defines and governs the behavior of the privileged resources included in a MIPS® processor implementation
- Volume IV-a describes the MIPS16e™ Application-Specific Extension to the MIPS64® Architecture. Beginning with Release 3 of the Architecture, microMIPS is the preferred solution for smaller code size.
- Volume IV-b describes the MDMX™ Application-Specific Extension to the MIPS64® Architecture and microMIPS64™. With Release 5 of the Architecture, MDMX is deprecated. MDMX and MSA can not be implemented at the same time.
- Volume IV-c describes the MIPS-3D® Application-Specific Extension to the MIPS® Architecture
- Volume IV-d describes the SmartMIPS® Application-Specific Extension to the MIPS32® Architecture and the microMIPS32™ Architecture and is not applicable to the MIPS64® document set nor the microMIPS64™ document set.
- Volume IV-e describes the MIPS® DSP Module to the MIPS® Architecture
- Volume IV-f describes the MIPS® MT Module to the MIPS® Architecture
- Volume IV-h describes the MIPS® MCU Application-Specific Extension to the MIPS® Architecture
- Volume IV-i describes the MIPS® Virtualization Module to the MIPS® Architecture
- Volume IV-j describes the MIPS® SIMD Architecture Module to the MIPS® Architecture

1.1 Typographical Conventions

This section describes the use of *italic*, **bold** and `courier` fonts in this book.

1.1.1 Italic Text

- is used for *emphasis*
- is used for *bits, fields, registers*, that are important from a software perspective (for instance, address bits used by software, and programmable fields and registers), and various *floating point instruction formats*, such as *S*, *D*, and *PS*
- is used for the memory access types, such as *cached* and *uncached*

1.1.2 Bold Text

- represents a term that is being **defined**
- is used for **bits** and **fields** that are important from a hardware perspective (for instance, **register** bits, which are not programmable but accessible only to hardware)
- is used for ranges of numbers; the range is indicated by an ellipsis. For instance, **5..1** indicates numbers 5 through 1
- is used to emphasize **UNPREDICTABLE** and **UNDEFINED** behavior, as defined below.

1.1.3 Courier Text

`Courier` fixed-width font is used for text that is displayed on the screen, and for examples of code and instruction pseudocode.

1.2 UNPREDICTABLE and UNDEFINED

The terms **UNPREDICTABLE** and **UNDEFINED** are used throughout this book to describe the behavior of the processor in certain cases. **UNDEFINED** behavior or operations can occur only as the result of executing instructions in a privileged mode (i.e., in Kernel Mode or Debug Mode, or with the CP0 usable bit set in the Status register). Unprivileged software can never cause **UNDEFINED** behavior or operations. Conversely, both privileged and unprivileged software can cause **UNPREDICTABLE** results or operations.

1.2.1 UNPREDICTABLE

UNPREDICTABLE results may vary from processor implementation to implementation, instruction to instruction, or as a function of time on the same implementation or instruction. Software can never depend on results that are **UNPREDICTABLE**. **UNPREDICTABLE** operations may cause a result to be generated or not. If a result is generated, it is **UNPREDICTABLE**. **UNPREDICTABLE** operations may cause arbitrary exceptions.

UNPREDICTABLE results or operations have several implementation restrictions:

- Implementations of operations generating **UNPREDICTABLE** results must not depend on any data source (memory or internal state) which is inaccessible in the current processor mode
- **UNPREDICTABLE** operations must not read, write, or modify the contents of memory or internal state which is inaccessible in the current processor mode. For example, **UNPREDICTABLE** operations executed in user mode must not access memory or internal state that is only accessible in Kernel Mode or Debug Mode or in another process

- **UNPREDICTABLE** operations must not halt or hang the processor

1.2.2 UNDEFINED

UNDEFINED operations or behavior may vary from processor implementation to implementation, instruction to instruction, or as a function of time on the same implementation or instruction. **UNDEFINED** operations or behavior may vary from nothing to creating an environment in which execution can no longer continue. **UNDEFINED** operations or behavior may cause data loss.

UNDEFINED operations or behavior has one implementation restriction:

- **UNDEFINED** operations or behavior must not cause the processor to hang (that is, enter a state from which there is no exit other than powering down the processor). The assertion of any of the reset signals must restore the processor to an operational state

1.2.3 UNSTABLE

UNSTABLE results or values may vary as a function of time on the same implementation or instruction. Unlike **UNPREDICTABLE** values, software may depend on the fact that a sampling of an **UNSTABLE** value results in a legal transient value that was correct at some point in time prior to the sampling.

UNSTABLE values have one implementation restriction:

- Implementations of operations generating **UNSTABLE** results must not depend on any data source (memory or internal state) which is inaccessible in the current processor mode

1.3 Special Symbols in Pseudocode Notation

In this book, algorithmic descriptions of an operation are described as pseudocode in a high-level language notation resembling Pascal. Special symbols used in the pseudocode notation are listed in [Table 1.1](#).

Table 1.1 Symbols Used in Instruction Operation Statements

Symbol	Meaning
\leftarrow	Assignment
$=, \neq$	Tests for equality and inequality
\parallel	Bit string concatenation
x^y	A y -bit string formed by y copies of the single-bit value x
$b\#n$	A constant value n in base b . For instance $10\#100$ represents the decimal value 100, $2\#100$ represents the binary value 100 (decimal 4), and $16\#100$ represents the hexadecimal value 100 (decimal 256). If the "b#" prefix is omitted, the default base is 10.
$0bn$	A constant value n in base 2. For instance $0b100$ represents the binary value 100 (decimal 4).
$0xn$	A constant value n in base 16. For instance $0x100$ represents the hexadecimal value 100 (decimal 256).
$x_{y..z}$	Selection of bits y through z of bit string x . Little-endian bit notation (rightmost bit is 0) is used. If y is less than z , this expression is an empty (zero length) bit string.
$+, -$	2's complement or floating point arithmetic: addition, subtraction

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning
$*, \times$	2's complement or floating point multiplication (both used for either)
div	2's complement integer division
mod	2's complement modulo
/	Floating point division
<	2's complement less-than comparison
>	2's complement greater-than comparison
\leq	2's complement less-than or equal comparison
\geq	2's complement greater-than or equal comparison
nor	Bitwise logical NOR
xor	Bitwise logical XOR
and	Bitwise logical AND
or	Bitwise logical OR
not	Bitwise inversion
&&	Logical (non-Bitwise) AND
<<	Logical Shift left (shift in zeros at right-hand-side)
>>	Logical Shift right (shift in zeros at left-hand-side)
GPRLen	The length in bits (32 or 64) of the CPU general-purpose registers
$GPR[x]$	CPU general-purpose register x . The content of $GPR[0]$ is always zero. In Release 2 of the Architecture, $GPR[x]$ is a short-hand notation for $SGPR[SRSCtl_{CSS}, x]$.
$SGPR[s, x]$	In Release 2 of the Architecture and subsequent releases, multiple copies of the CPU general-purpose registers may be implemented. $SGPR[s, x]$ refers to GPR set s , register x .
$FPR[x]$	Floating Point operand register x
$FCC[CC]$	Floating Point condition code CC . $FCC[0]$ has the same value as $COC[1]$.
$FPR[x]$	Floating Point (Coprocessor unit 1), general register x
$CPR[z, x, s]$	Coprocessor unit z , general register x , select s
CP2CPR[x]	Coprocessor unit 2, general register x
$CCR[z, x]$	Coprocessor unit z , control register x
CP2CCR[x]	Coprocessor unit 2, control register x
$COC[z]$	Coprocessor unit z condition signal
$Xlat[x]$	Translation of the MIPS16e GPR number x into the corresponding 32-bit GPR number
BigEndianMem	Endian mode as configured at chip reset (0 \rightarrow Little-Endian, 1 \rightarrow Big-Endian). Specifies the endianness of the memory interface (see LoadMemory and StoreMemory pseudocode function descriptions), and the endianness of Kernel and Supervisor mode execution.
BigEndianCPU	The endianness for load and store instructions (0 \rightarrow Little-Endian, 1 \rightarrow Big-Endian). In User mode, this endianness may be switched by setting the RE bit in the <i>Status</i> register. Thus, BigEndianCPU may be computed as (BigEndianMem XOR ReverseEndian).
ReverseEndian	Signal to reverse the endianness of load and store instructions. This feature is available in User mode only, and is implemented by setting the RE bit of the <i>Status</i> register. Thus, ReverseEndian may be computed as (SR_{RE} and User mode).

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning						
<i>LLbit</i>	Bit of virtual state used to specify operation for instructions that provide atomic read-modify-write. <i>LLbit</i> is set when a linked load occurs and is tested by the conditional store. It is cleared, during other CPU operation, when a store to the location would no longer be atomic. In particular, it is cleared by exception return instructions.						
I , I+n , I-n :	This occurs as a prefix to <i>Operation</i> description lines and functions as a label. It indicates the instruction time during which the pseudocode appears to “execute.” Unless otherwise indicated, all effects of the current instruction appear to occur during the instruction time of the current instruction. No label is equivalent to a time label of I . Sometimes effects of an instruction appear to occur either earlier or later — that is, during the instruction time of another instruction. When this happens, the instruction operation is written in sections labeled with the instruction time, relative to the current instruction I , in which the effect of that pseudocode appears to occur. For example, an instruction may have a result that is not available until after the next instruction. Such an instruction has the portion of the instruction operation description that writes the result register in a section labeled I+1 . The effect of pseudocode statements for the current instruction labelled I+1 appears to occur “at the same time” as the effect of pseudocode statements labeled I for the following instruction. Within one pseudocode sequence, the effects of the statements take place in order. However, between sequences of statements for different instructions that occur “at the same time,” there is no defined order. Programs must not depend on a particular order of evaluation between such sections.						
PC	The <i>Program Counter</i> value. During the instruction time of an instruction, this is the address of the instruction word. The address of the instruction that occurs during the next instruction time is determined by assigning a value to <i>PC</i> during an instruction time. If no value is assigned to <i>PC</i> during an instruction time by any pseudocode statement, it is automatically incremented by either 2 (in the case of a 16-bit MIPS16e instruction) or 4 before the next instruction time. A taken branch assigns the target address to the <i>PC</i> during the instruction time of the instruction in the branch delay slot. In the MIPS Architecture, the PC value is only visible indirectly, such as when the processor stores the restart address into a GPR on a jump-and-link or branch-and-link instruction, or into a Coprocessor 0 register on an exception. The PC value contains a full 64-bit address all of which are significant during a memory reference.						
ISA Mode	In processors that implement the MIPS16e Application Specific Extension or the microMIPS base architectures, the <i>ISA Mode</i> is a single-bit register that determines in which mode the processor is executing, as follows: <table border="1" data-bbox="597 1251 1266 1398"> <thead> <tr> <th>Encoding</th><th>Meaning</th></tr> </thead> <tbody> <tr> <td>0</td><td>The processor is executing 32-bit MIPS instructions</td></tr> <tr> <td>1</td><td>The processor is executing MIPS16e or microMIPS instructions</td></tr> </tbody> </table> <p>In the MIPS Architecture, the ISA Mode value is only visible indirectly, such as when the processor stores a combined value of the upper bits of PC and the ISA Mode into a GPR on a jump-and-link or branch-and-link instruction, or into a Coprocessor 0 register on an exception.</p>	Encoding	Meaning	0	The processor is executing 32-bit MIPS instructions	1	The processor is executing MIPS16e or microMIPS instructions
Encoding	Meaning						
0	The processor is executing 32-bit MIPS instructions						
1	The processor is executing MIPS16e or microMIPS instructions						
PABITS	The number of physical address bits implemented is represented by the symbol PABITS. As such, if 36 physical address bits were implemented, the size of the physical address space would be $2^{\text{PABITS}} = 2^{36}$ bytes.						
SEGBITS	The number of virtual address bits implemented in a segment of the address space is represented by the symbol SEGBITS. As such, if 40 virtual address bits are implemented in a segment, the size of the segment is $2^{\text{SEGBITS}} = 2^{40}$ bytes.						

Table 1.1 Symbols Used in Instruction Operation Statements (Continued)

Symbol	Meaning
FP32RegistersMode	Indicates whether the FPU has 32-bit or 64-bit floating point registers (FPRs). It is optional if the FPU has 32 64-bit FPRs in which 64-bit data types are stored in any FPR. microMIPS64 implementations have a compatibility mode in which the processor references the FPRs as if it were a microMIPS32 implementation. In such a case FP32RegisterMode is computed from the FR bit in the <i>Status</i> register. If this bit is a 0, the processor operates as if it had 32 32-bit FPRs. If this bit is a 1, the processor operates with 32 64-bit FPRs. The value of FP32RegistersMode is computed from the FR bit in the <i>Status</i> register.
InstructionInBranchDelaySlot	Indicates whether the instruction at the Program Counter address was executed in the delay slot of a branch or jump. This condition reflects the <i>dynamic</i> state of the instruction, not the <i>static</i> state. That is, the value is false if a branch or jump occurs to an instruction whose PC immediately follows a branch or jump, but which is not executed in the delay slot of a branch or jump.
SignalException(exception, argument)	Causes an exception to be signaled, using the exception parameter as the type of exception and the argument parameter as an exception-specific argument). Control does not return from this pseudocode function—the exception is signaled at the point of the call.

1.4 For More Information

Various MIPS RISC processor manuals and additional information about MIPS products can be found at the MIPS URL: <http://www.mips.com>

For comments or questions on the MIPS64® Architecture or this document, send Email to support@mips.com.

The MIPS Architecture: An Introduction

2.1 MIPS Instruction Set Overview

2.1.1 Historical Perspective

The MIPS® Instruction Set Architecture (ISA) has evolved over time from the original MIPS I™ ISA, through the MIPS V™ ISA, to the current MIPS32®, MIPS64® and microMIPS™ Architectures. As the ISA evolved, all extensions have been backward compatible with previous versions of the ISA. In the MIPS III™ level of the ISA, 64-bit integers and addresses were added to the instruction set. The MIPS IV™ and MIPS V™ levels of the ISA added improved floating point operations, as well as a set of instructions intended to improve the efficiency of generated code and of data movement. Because of the strict backward-compatible requirement of the ISA, such changes were unavailable to 32-bit implementations of the ISA which were, by definition, MIPS I™ or MIPS II™ implementations.

While the user-mode ISA was always backward compatible, the privileged environment was allowed to change on a per-implementation basis. As a result, the R3000® privileged environment was different from the R4000® privileged environment, and subsequent implementations, while similar to the R4000 privileged environment, included subtle differences. Because the privileged environment was never part of the MIPS ISA, an implementation had the flexibility to make changes to suit that particular implementation. Unfortunately, this required kernel software changes to every operating system or kernel environment on which that implementation was intended to run.

Many of the original MIPS implementations were targeted at computer-like applications such as workstations and servers. In recent years MIPS implementations have had significant success in embedded applications. Today, most of the MIPS parts that are shipped go into some sort of embedded application. Such applications tend to have different trade-offs than computer-like applications including a focus on cost of implementation, and performance as a function of cost and power.

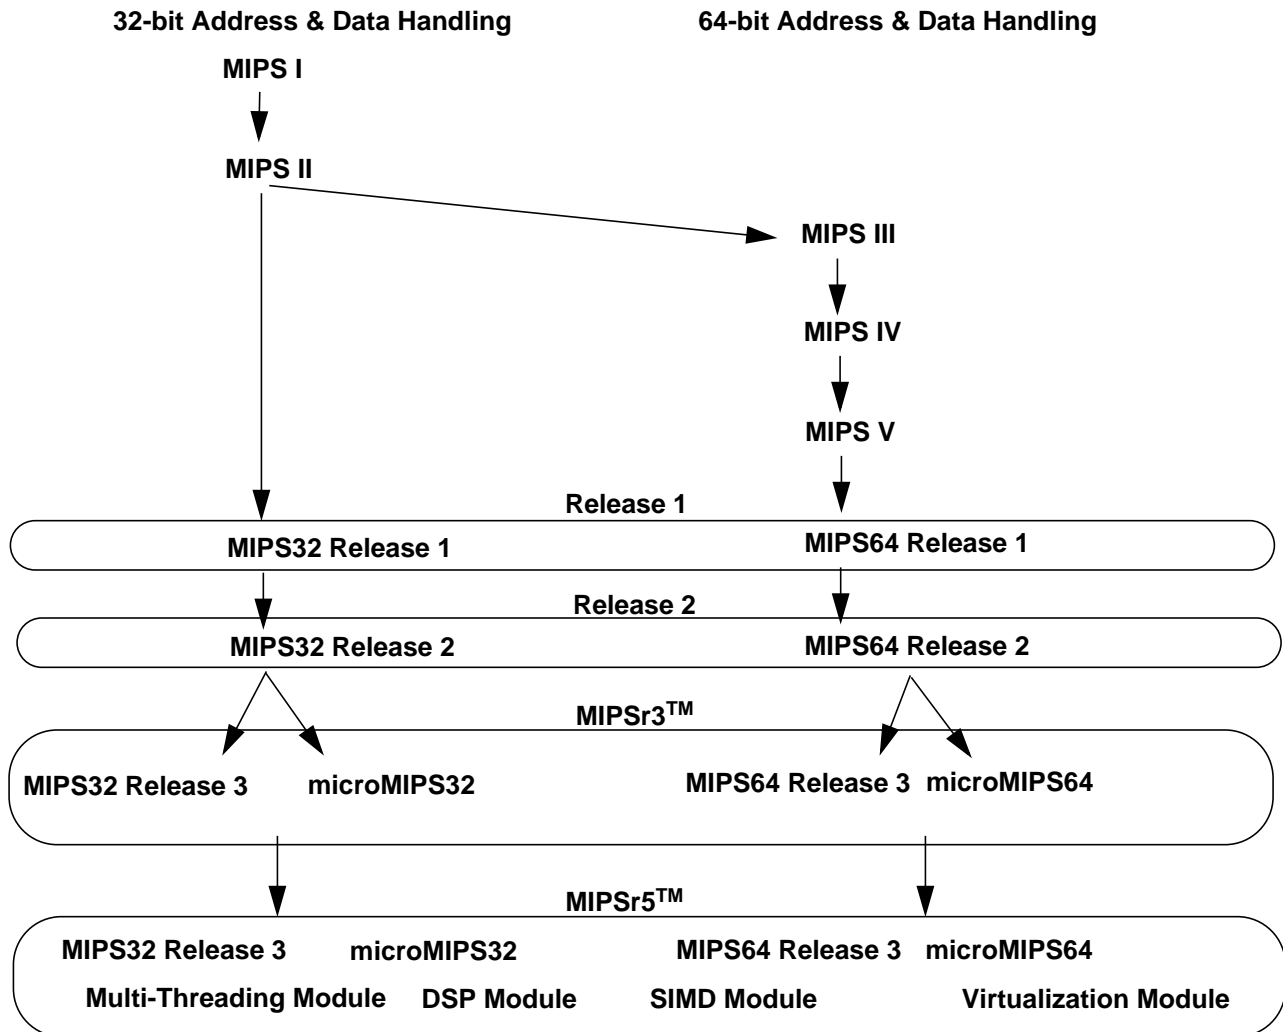
The MIPS32 and MIPS64 Architectures are intended to address the need for a high-performance but cost-sensitive MIPS instruction set. The MIPS32 Architecture is based on the MIPS II ISA, adding selected instructions from MIPS III, MIPS IV, and MIPS V to improve the efficiency of generated code and of data movement. The MIPS64 Architecture is based on the MIPS V ISA and is backward compatible with the MIPS32 Architecture. Both the MIPS32 and MIPS64 Architectures bring the privileged environment into the Architecture definition to address the needs of operating systems and other kernel software. Both also include provision for adding optional components - Modules of the base architecture, MIPS Application Specific Extensions (ASEs), User Defined Instructions (UDIs), and custom coprocessors to address the specific needs of particular markets.

The MIPS32 and MIPS64 Architectures provide a substantial cost/performance advantage over microprocessor implementations based on traditional architectures. This advantage is a result of improvements made in several contiguous disciplines: VLSI process technology, CPU organization, system-level architecture, and operating system and compiler design.

The microMIPS32 and microMIPS64 Architectures deliver the same functionality of MIPS32 and MIPS64 with the additional benefit of smaller codesizes. The microMIPS architectures are supersets of MIPS32/MIPS64 architectures, with almost the same sets of 32-bit sized instructions and additional 16-bit instructions to help with codesize. microMIPS is especially compelling for systems in which the cost of memory dominate the entire bill of materials cost.

Unlike the earlier versions of the architectures, microMIPS supplies assembler-source code compatibility with its predecessors instead of binary compatibility.

Figure 2.1 MIPS Architectures



2.1.2 Architectural Evolution

The evolution of an architecture is a dynamic process that takes into account both the need to provide a stable platform for implementations, as well as new market and application areas that demand new capabilities. Enhancements to an architecture are appropriate when they:

- are applicable to a wide market
- provide long-term benefit
- maintain architectural scalability

- are standardized to prevent fragmentation
- are a superset of the existing architecture

The MIPS Architecture community constantly evaluates suggestions for architectural changes and enhancements against these criteria. New releases of the architecture, while infrequent, are made at appropriate points, following these criteria. At present, there are three releases of the MIPS Architecture: Release 1 (the original version of the MIPS64 Architecture) ; Release 2 which was added in 2002 and Release 3 (called MIPSr3™) which was added in 2010.

2.1.2.1 Release 2 of the MIPS64 Architecture

Enhancements included in Release 2 of the MIPS64 Architecture are:

- **Vectored interrupts:** This enhancement provides the ability to vector interrupts directly to a handler for that interrupt. Vectored interrupts are an option in Release 2 implementations and the presence of that option is denoted by the Config3_{VInt} bit.
- **Support for an external interrupt controller:** This enhancement reconfigures the on-core interrupt logic to take full advantage of an external interrupt controller. This support is an option in Release 2 implementations and the presence of that option is denoted by the Config3_{EIC} bit.
- **Programmable exception vector base:** This enhancement allows the base address of the exception vectors to be moved for exceptions that occur when Status_{BEV} is 0. Doing so allows multi-processor systems to have separate exception vectors for each processor, and allows any system to place the exception vectors in memory that is appropriate to the system environment. This enhancement is required in a Release 2 implementation.
- **Atomic interrupt enable/disable:** Two instructions have been added to atomically enable or disable interrupts, and return the previous value of the *Status* register. These instructions are required in a Release 2 implementation.
- **The ability to disable the *Count* register for highly power-sensitive applications.** This enhancement is required in a Release 2 implementation.
- **GPR shadow registers:** This addition provides the addition of GPR shadow registers and the ability to bind these registers to a vectored interrupt or exception. Shadow registers are an option in Release 2 implementations and the presence of that option is denoted by a non-zero value in SRSCtl_{HSS}. While shadow registers are most useful when either vectored interrupts or support for an external interrupt controller is also implemented, neither is required.
- **Field, Rotate and Shuffle instructions:** These instructions add additional capability in processing bit fields in registers. These instructions are required in a Release 2 implementation.
- **Explicit hazard management:** This enhancement provides a set of instructions to explicitly manage hazards, in place of the cycle-based SSNOP method of dealing with hazards. These instructions are required in a Release 2 implementation.
- **Access to a new class of hardware registers and state from an unprivileged mode.** This enhancement is required in a Release 2 implementation.
- **Coprocessor 0 Register changes:** These changes add or modify CP0 registers to indicate the existence of new and optional state, provide L2 and L3 cache identification, add trigger bits to the Watch registers, and add support for 64-bit performance counter count registers. This enhancement is required in a Release 2 implementation.

- Support for 64-bit coprocessors with 32-bit CPUs: These changes allow a 64-bit coprocessor (including an FPU) to be attached to a 32-bit CPU. This enhancement is optional in a Release 2 implementation.
- New Support for Virtual and Physical Memory: These changes provide support for a 1KByte page size, and the ability to support physical addresses larger than 36 bits. Both changes are optional in Release 2 implementations, and support is denoted by Config3_{SP} (for 1KB page support) and Config3_{LPA} (for larger physical address support).

2.1.2.2 Releases 2.5+ of the MIPS64 Architecture

Some optional features were added after Revision 2.5:

- TLB pages larger than 256MB are supported. This feature allows large regions to be mapped with fewer TLB entries, especially within devices with very large memory systems.
- Support for a MMU with more than 64 TLB entries. This feature aids in reducing the frequency of TLB misses.
- Scratch registers within Coprocessor0 for kernel mode software. This feature aids in quicker exception handling by not requiring the saving of usermode registers onto the stack before kernelmode software uses those registers.
- A MMU configuration which supports both larger set-associative TLBs and variable page-sizes. This feature aids in reducing the frequency of TLB misses.
- The CDMM memory scheme for the placement of small I/O devices into the physical address space. This scheme allows for efficient placement of such I/O devices into a small memory region.
- An EIC interrupt mode where the EIC controller supplies a 16-bit interrupt vector. This allows different interrupts to share code.
- The PAUSE instruction to deallocate a (virtual) processor when arbitration for a lock doesn't succeed. This allows for lower power consumption as well as lower snoop traffic when multiple (virtual) processors are arbitrating for a lock.
- More flavors of memory barriers that are available through stype field of the SYNC instruction. The newer memory barriers attempt to minimize the amount of pipeline stalls while doing memory synchronization operations.

2.1.2.3 MIPSr3™ Architecture

MIPSr3™ is a family of architectures which includes Release 3.0 of the MIPS64 Architecture as well as the first release of the microMIPS64 architecture.

Enhancements included in MIPSr3™ Architecture are:

- The microMIPS™ instruction set.
 - This instruction set contains both 16-bit and 32-bit sized instructions.
 - This mixed size ISA has all of the functionality of MIPS64 while also delivering smaller code sizes.
 - microMIPS is assembler source code compatible with MIPS64.
 - microMIPS replaces the MIPS16e™ ASE.

- microMIPS is an additional base instruction set architecture that is supported along with MIPS64.
- A device can implement either base ISA or both. The ISA field of *Config3* denotes which ISA is implemented.
- A device can implement any other Module/ASE with either base architecture.¹
- microMIPS shares the same privileged resource architecture with MIPS64.
- Branch Likely instructions are not supported in the microMIPS hardware architecture. Instead the microMIPS toolchain replaces these instructions with equivalent code sequences.
- A more flexible version of the Context Register that can point to any power-of-two sized data structure. This optional feature is denoted by CTXTC field of *Config3*.
- Additional protection bits in the TLB entries that allow for non-executable and write-only virtual pages. This optional feature is denoted by RXI field of *Config3*.
- A more programmable virtual address space map without fixed cache-ability and map-ability attributes is introduced as an optional feature. This allows the implementations to decide how large/small uncached/unmapped segments need to be. These capabilities are implemented through the Segmentation Control registers. This optional feature is denoted by SC field of *Config3*.
- Along with programmable virtual address map, it is possible to create separate user-mode & kernel-mode views of segments. This allows a larger kernel virtual address space to be defined. To access both this larger kernel address space and the overlapping user-space, additional load/store instructions are introduced. These new optional instructions are denoted by EVA field of *Config5*.
- Support for certain IEEE-754-2008 FPU behaviors (as opposed to behaviors of the older IEEE-754-1985 standard) is now defined. These behaviors are indicated by the Has2008 field of the FIR register within the FPU and FCSR bits ABS2008 or NAN2008.²
- Optional TLB invalidate instructions are introduced. These are necessary with Segmentation Control as it is now possible to create a virtual address map without unmapped segments.

2.1.2.4 MIPSr5™ Architecture

MIPSr5™ is a family of architectures (MIPS32, MIPS64, microMIPS32 and microMIPS64) and adds these capabilities:

- The Multi-threading module is now an optional component of all of these base architectures. Previously the MT ASE was sold as a separate architecture product.
- The DSP module is now an optional component of all of these base architectures. Previously the DSP ASE was sold as a separate architecture product.
- The Virtualization module is introduced for all of these base architectures.

1. Except for MIPS16e.

2. At this time the MIPS32 and MIPS64 architectures provide no feature supporting IEEE-754-2008 fused multiply add without intermediate rounding.

- The MIPS SIMD Architecture (MSA) module is introduced for all of these base architectures.

In addition, these changes are made:

- The MDMX ASE is formally deprecated. The equivalent functionality is covered by the MSA module.
- The 64-bit versions of the DSP ASE are formally deprecated. The equivalent functionality is covered by the MSA module.
- As of Release 5 of the Architecture, if floating point is implemented then FR=1 is required. I.e. the 64-bit FPU, with the FR=1 64-bit FPU register model, is required. The FR=0 32-bit FPU register model continues to be required.³

2.1.3 Architectural Changes Relative to the MIPS I through MIPS V Architectures

In addition to the MIPS Architecture described in this document set, the following changes were made to the architecture relative to the earlier MIPS RISC Architecture Specification, which describes the MIPS I through MIPS V Architectures.

- The MIPS IV ISA added a restriction to the load and store instructions which have natural alignment requirements (all but load and store byte and load and store left and right) in which the base register used by the instruction must also be naturally aligned (the restriction expressed in the MIPS RISC Architecture Specification is that the offset be aligned, but the implication is that the base register is also aligned, and this is more consistent with the indexed load/store instructions which have no offset field). The restriction that the base register be naturally-aligned is eliminated by the MIPS64 Architecture, leaving the restriction that the effective address be naturally-aligned.
- Early MIPS implementations required two instructions separating a *MFLO* or *MFHI* from the next integer multiply or divide operation. This hazard was eliminated in the MIPS IV ISA, although the MIPS RISC Architecture Specification does not clearly explain this fact. The MIPS64 Architecture explicitly eliminates this hazard and requires that the hi and lo registers be fully interlocked in hardware for all integer multiply and divide instructions (including, but not limited to, the *MADD*, *MADDU*, *MSUB*, *MSUBU*, and *MUL* instructions introduced in this specification).
- The Implementation and Programming Notes included in the instruction descriptions for the *madd*, *maddu*, *msub*, *msubu*, and *mul* instructions should also be applied to all integer multiply and divide instructions in the MIPS RISC Architecture Specification.

2.2 Compliance and Subsetting

To be compliant with the microMIPS64 Architecture, designs must implement a set of required features, as described in this document set. To allow flexibility in implementations, the microMIPS64 Architecture does provide subsetting rules. An implementation that follows these rules is compliant with the microMIPS64 Architecture as long as it adheres strictly to the rules, and fully implements the remaining instructions. Supersetting of the microMIPS64 Architecture is only allowed by adding functions to the *SPECIAL2* opcode, by adding control for co-processors via the *COP2*, *LWC2*, *SWC2*, *LDC2*, and/or *SDC2*, or via the addition of approved Application Specific Extensions.

-
3. Release 5 of the Architecture makes the FPU requirements consistent between MIPS32 and MIPS64. Prior to Release 5 MIPS64 requires FR=0 and FR=1, whereas MIPS32 requires FR=0 but FR=1 is optional. Release 5 requires FR=0 and FR=1 in all implementations of floating point, although floating point overall remains optional.

Note: The use of COP3 as a customizable coprocessor has been removed in the Release 2 of the MIPS64 architecture. The use of the COP3 is now reserved for the future extension of the architecture.

The instruction set subsetting rules are as follows:

- All non-privileged (do not need access to Coprocessor 0) CPU (non-FPU) instructions must be implemented - no subsetting of these are allowed.
- The FPU and related support instructions, including the MOVF and MOVT CPU instructions, may be omitted. Software may determine if an FPU is implemented by checking the state of the FP bit in the *Config1* CP0 register. If the FPU is implemented, the paired single (PS) format is optional. Software may determine which FPU data types are implemented by checking the appropriate bit in the *FIR* CP1 register. The following allowable FPU subsets are compliant with the MIPS64 architecture:
 - No FPU
 - FPU with S, D, W, and L formats and all supporting instructions
 - FPU with S, D, PS, W, and L formats and all supporting instructions
 - As of Release 5 of the Architecture, if floating point is implemented then FR=1 is required. I.e. the 64-bit FPU, with the FR=1 64-bit FPU register model, is required. The FR=0 32-bit FPU register model continues to be required.
- Coprocessor 2 is optional and may be omitted. Software may determine if Coprocessor 2 is implemented by checking the state of the C2 bit in the *Config1* CP0 register. If Coprocessor 2 is implemented, the Coprocessor 2 interface instructions (BC2, CFC2, COP2, CTC2, DMFC2, DMTC2, LDC2, LWC2, MFC2, MTC2, SDC2, and SWC2) may be omitted on an instruction-by-instruction basis.
- Implementation of the full 64-bit address space is optional. The processor may implement 64-bit data and operations with a 32-bit only address space. In this case, the MMU acts as if 64-bit addressing is always disabled. Software may determine if the processor implements a 32-bit or 64-bit address space by checking the AT field in the *Config* CP0 register.
- The EVA load/store instructions (LWE, LHE, LBE, LBUE, LHUE, SWE, SHE, SBE) are optional.
- Supervisor Mode is optional. If Supervisor Mode is not implemented, bit 3 of the *Status* register must be ignored on write and read as zero.
- The standard TLB-based memory management unit may be replaced with:
 - a simpler MMU (e.g., a Fixed Mapping MMU or a Block Address Translation MMU or a Base-Bounds MMU).
 - The Dual TLB MMU - (e.g. FTLB and VTLB MMU described in the *Alternative MMU Organizations* Appendix of Volume III)

If this is done, the rest of the interface to the Privileged Resource Architecture must be preserved. Software may determine the type of the MMU by checking the MT field in the *Config* CP0 register.

- The Caches are optional. The *Config1_{DL}* and *Config1_{IL}* fields denote whether the first level caches are present or not.

- The Privileged Resource Architecture includes several implementation options and may be subsetting in accordance with those options. An incomplete list of these options include:
 - Interrupt Modes
 - Shadow Register Sets
 - Common Device Memory Map
 - Parity/ECC support
 - UserLocal register
 - ContextConfig register
 - PageGrain register
 - Config1-4 registers
 - Performance Counter, WatchPoint and Trace Registers
 - Cache control/diagnostic registers
 - Kernelmode scratch registers
- Instruction, CP0 Register, and CP1 Control Register fields that are marked “Reserved” or shown as “0” in the description of that field are reserved for future use by the architecture and are not available to implementations. Implementations may only use those fields that are explicitly reserved for implementation dependent use.
- Supported Modules/ASEs are optional and may be subsetting out. In most cases, software may determine if a supported Module/ASE is implemented by checking the appropriate bit in the *Config1* or *Config3* or *Config4* CP0 register. If they are implemented, they must implement the entire ISA applicable to the component, or implement subsets that are approved by the Module/ASE specifications.
- EJTAG is optional and may be subsetting out. If it is implemented, it must implement only those subsets that are approved by the EJTAG specification.
- If any instruction is subsetting out based on the rules above, an attempt to execute that instruction must cause the appropriate exception (typically Reserved Instruction or Coprocessor Unusable).
- In MIPSr3 (also called Release 3), there are two architecture branches (MIPS32/64 and microMIPS32/64). A single device is allowed to implement both architecture branches. The Privileged Resource Architecture (COP0) registers do not mode-switch in width (32-bit vs. 64-bit). For this reason, if a device implements both architecture branches, the address/data widths must be consistent. If a device implements MIPS64 and also implements microMIPS, it must implement microMIPS64 not just microMIPS32. Similarly, If a device implements microMIPS64 and also implements MIPS32/64, it must implement MIPS64 not just MIPS32.
- If both of the architecture branches are implemented (MIPS32/64 and microMIPS32/64) or if MIPS16e is implemented then the JALX instructions are required. If only one branch of the architecture family and MIPS16e is not implemented then the JALX instruction is not implemented. That is, the JALX instruction is required if and only if when ISA mode-switching is possible.

- MIPSr5™ (also called Release 5) includes a number of features. Some are optional; some are required. Release 5 features, whether optional or required, must be consistent. If any feature that is introduced by Release 5 is implemented, i.e. which is described as part of Release 5 and not any earlier release, then all other features must be implemented in a manner consistent with Release 5. For example: the VZ and MSA features are introduced by Release 5 but are optional, whereas the FR=1 64-bit FPU register model was optional when introduced earlier, but is now required by Release 5 if any FPU is implemented. If any or all of VZ or MSA are implemented, then Release 5 is implied, and then if an FPU is implemented, it must implement the FR=1 64-bit FPU register model.

2.3 Components of the MIPS Architecture

2.3.1 MIPS Instruction Set Architecture (ISA)

The microMIPS32 and microMIPS64 Instruction Set Architectures define a compatible family of instructions dealing with 32-bit data and 64-bit data (respectively) within the framework of the overall MIPS Architectures. Included in the ISA are all instructions, both privileged and unprivileged, by which the programmer interfaces with the processor. The ISA guarantees object code compatibility for unprivileged and, often, privileged programs executing on any microMIPS32 or microMIPS64 processor; all instructions in the microMIPS64 ISA are backward compatible with those instructions in the microMIPS32 ISA. Using conditional compilation or assembly language macros, it is often possible to write privileged programs that run on both MIPS32 and MIPS64 implementations.

2.3.2 MIPS Privileged Resource Architecture (PRA)

The microMIPS32 and microMIPS64 Privileged Resource Architecture defines a set of environments and capabilities on which the ISA operates. The effects of some components of the PRA are visible to unprivileged programs; for instance, the virtual memory layout. Many other components are visible only to privileged programs and the operating system. The PRA provides the mechanisms necessary to manage the resources of the processor: virtual memory, caches, exceptions, user contexts, etc.

2.3.3 MIPS Modules and Application Specific Extensions (ASEs)

The microMIPS32 and microMIPS64 Architectures provide support for optional components - known as either Modules or application specific extensions. As optional extensions to the base architecture, the Modules/ASEs do not burden every implementation of the architecture with instructions or capability that are not needed in a particular market. An ASE/Module can be used with the appropriate ISA and PRA to meet the needs of a specific application or an entire class of applications.

2.3.4 MIPS User Defined Instructions (UDIs)

In addition to support for Modules/ASEs as described above, the MIPS32 and MIPS64 Architectures define specific instructions for the use of each implementation. The *Special2* instruction function fields and Coprocessor 2 are reserved for capability defined by each implementation.

2.4 Architecture Versus Implementation

When describing the characteristics of MIPS processors, *architecture* must be distinguished from the hardware *implementation of that architecture*.

- **Architecture** refers to the instruction set, registers and other state, the exception model, memory management, virtual and physical address layout, and other features that all hardware executes.

- **Implementation** refers to the way in which specific processors apply the architecture.

Here are two examples:

1. A floating point unit (FPU) is an optional part of the microMIPS64 Architecture. A compatible implementation of the FPU may have different pipeline lengths, different hardware algorithms for performing multiplication or division, etc.
2. Most MIPS processors have caches; however, these caches are not implemented in the same manner in all MIPS processors. Some processors implement physically-indexed, physically tagged caches. Other implement virtually-indexed, physically-tagged caches. Still other processor implement more than one level of cache.

The microMIPS64 architecture is decoupled from specific hardware implementations, leaving microprocessor designers free to create their own hardware designs within the framework of the architectural definition.

2.5 Relationship between the MIPSr3 Architectures

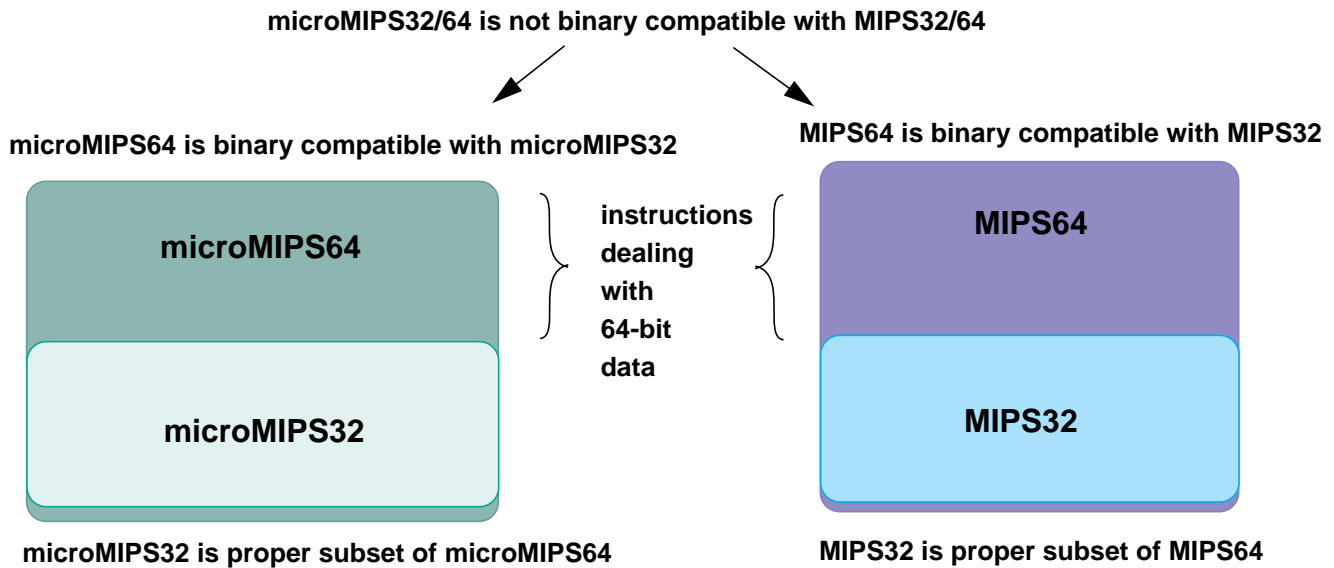
The MIPS Architectures evolved as a compromise between software and hardware resources. The MIPS has a family of related architectures. Within each “branch of the family”, the architecture guarantees object-code compatibility for User-Mode programs executed on any MIPS processor.

MIPS32 and MIPS64 form one branch of the architecture family. In User Mode MIPS64 processors are backward-compatible with their MIPS32 predecessors. As such, the MIPS32 Architecture is a strict subset of the MIPS64 Architecture.

Similarly, microMIPS32 and microMIPS64 form another branch of the architecture family. In User Mode microMIPS64 processors are backward-compatible with their microMIPS predecessors. As such, the microMIPS Architecture is a strict subset of the MIPS64 Architecture.

The relationship between the binary representations of the architectures is shown in [Figure 2.2](#).

Figure 2.2 Relationship of the Binary Representations of MIPSr3 Architectures




As of 2010, there are two branches of the architecture family - the MIPS32/64 branch and the microMIPS32/64 branch. For these two branches, some levels of compatibility are available:

1. The microMIPS32/64 branch supplies a superset of the functionality that is available from the MIPS32/64 branch. The additional functionality that the microMIPS branch delivers is smaller code size.
2. It is allowed for implementations to implement both branches of the architecture family for compatibility reasons. For such implementations, the architectures define methods of switching from one instruction set to the other. This allows one binary program to use both instruction sets or call a library that is using the other instruction set.
3. At the assembler source code level, the two architecture branches are fully compatible. That is, all of the MIPS32/64 assembler instruction mnemonics and directives are fully usable and understood by the microMIPS32/64 toolchains.

The relationships between the assembler source-code representations of the architectures is shown in [Figure 2.3](#).

Figure 2.3 Relationships of the Assembler Source Code Representations of the MIPSr3 Architectures



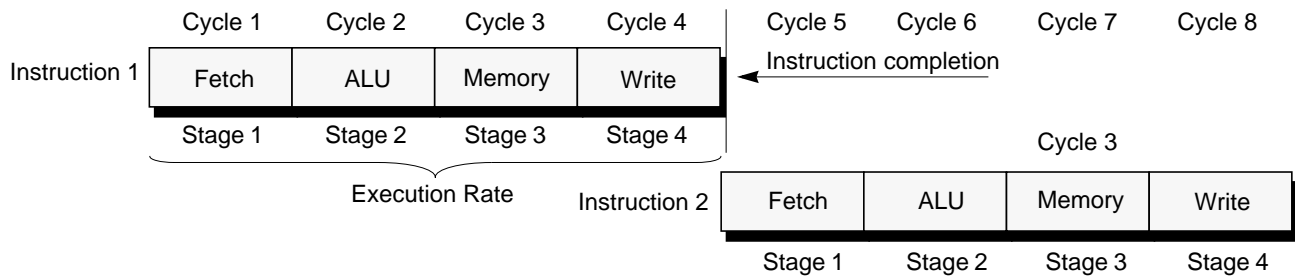
2.6 Pipeline Architecture

This section describes the basic pipeline architecture, along with two types of improvements: superpipelines and superscalar pipelines. (Pipelining and multiple issuing are not defined by the ISA, but are implementation dependent.)

2.6.1 Pipeline Stages and Execution Rates

MIPS processors all use some variation of a pipeline in their architecture. A pipeline is divided into the following discrete parts, or **stages**, shown in [Figure 2.4](#):

- Fetch
- Arithmetic operation
- Memory access
- Write back

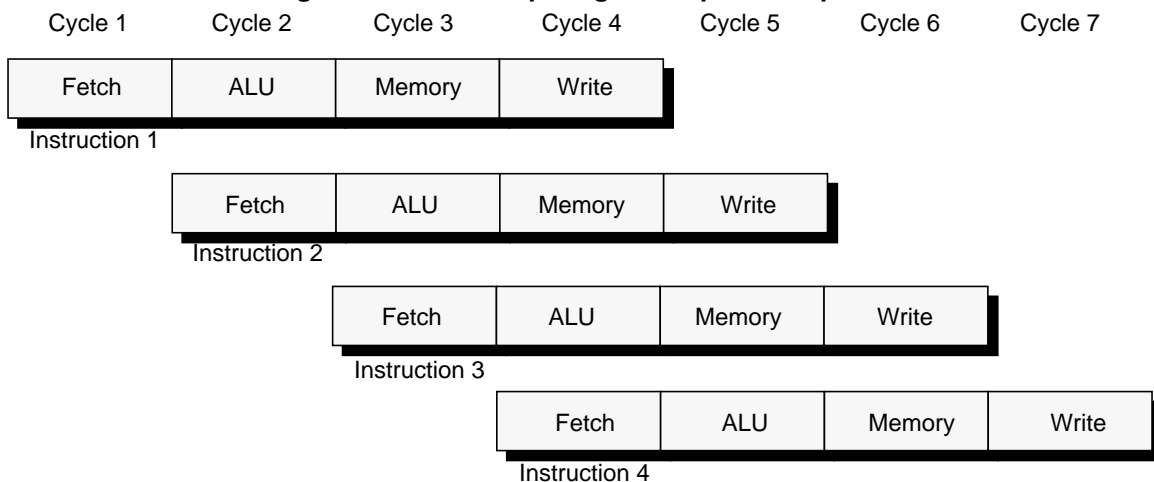
Figure 2.4 One-Deep Single-Completion Instruction Pipeline

In the example shown in [Figure 2.4](#), each stage takes one processor clock cycle to complete. Thus it takes four clock cycles (ignoring delays or stalls) for the instruction to complete. In this example, the **execution rate** of the pipeline is one instruction every four clock cycles. Conversely, because only a single execution can be fetched before completion, only one stage is active at any time.

2.6.2 Parallel Pipeline

[Figure 2.5](#) illustrates a remedy for the **latency** (the time it takes to execute an instruction) inherent in the pipeline shown in [Figure 2.4](#).

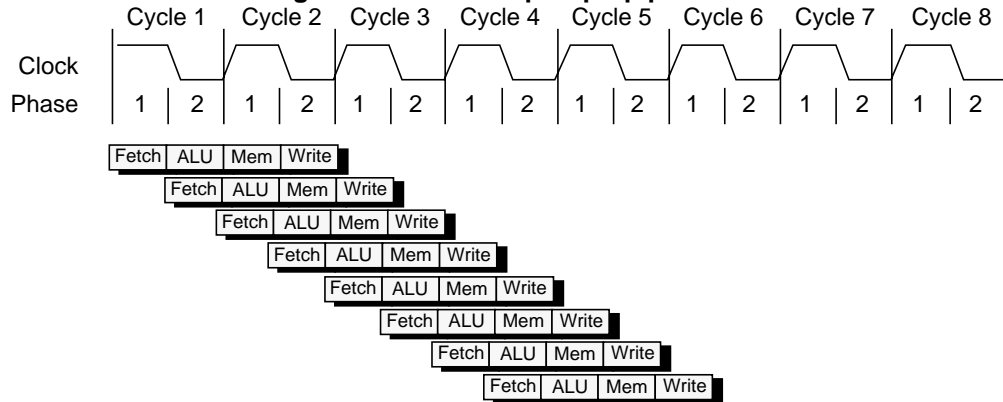
Instead of waiting for an instruction to be completed before the next instruction can be fetched (four clock cycles), a new instruction is fetched each clock cycle. There are four stages to the pipeline so the four instructions can be executed simultaneously, one at each stage of the pipeline. It still takes four clock cycles for the first instruction to be completed; however, in this theoretical example, a new instruction is completed every clock cycle thereafter. Instructions in [Figure 2.5](#) are executed at a rate four times that of the pipeline shown in [Figure 2.4](#).

Figure 2.5 Four-Deep Single-Completion Pipeline

2.6.3 Superpipeline

[Figure 2.6](#) shows a **superpipelined** architecture. Each stage is designed to take only a fraction of an external clock cycle—in this case, half a clock. Effectively, each stage is divided into more than one **substage**. Therefore more than one instruction can be completed each cycle.

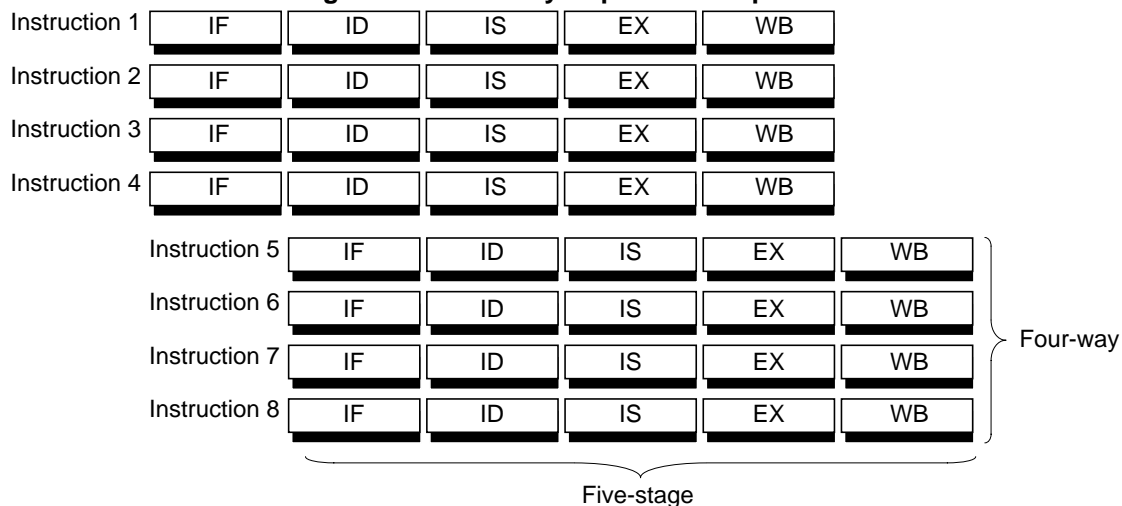
Figure 2.6 Four-Deep Superpipeline



2.6.4 Superscalar Pipeline

A **superscalar** architecture also allows more than one instruction to be completed each clock cycle. [Figure 2.7](#) shows a four-way, five-stage superscalar pipeline.

Figure 2.7 Four-Way Superscalar Pipeline



IF = instruction fetch
 ID = instruction decode and dependency
 IS = instruction issue
 EX = execution
 WB = write back

2.7 Load/Store Architecture

Generally, it takes longer to perform operations in memory than it does to perform them in on-chip registers. This is because of the difference in time it takes to access a register (fast) and main memory (slower).

To eliminate the longer access time, or **latency**, of in-memory operations, MIPS processors use a **load/store** design. The processor has many registers on chip, and all operations are performed on operands held in these processor registers. Main memory is accessed only through load and store instructions. This has several benefits:

- Reducing the number of memory accesses, easing memory bandwidth requirements
- Simplifying the instruction set
- Making it easier for compilers to optimize register allocation

2.8 Programming Model

This section describes the following aspects of the programming model:

- [CPU Data Formats](#)
- [Coproductors \(CP0-CP3\)](#)
- [CPU Registers](#)
- [FPU Data Formats](#)
- [Byte Ordering and Endianness](#)
- [Memory Access Types](#)

2.8.1 CPU Data Formats

The CPU defines the following data formats:

- Bit (*b*)
- Byte (8 bits, *B*)
- Halfword (16 bits, *H*)
- Word (32 bits, *W*)
- Doubleword (64 bits, *D*)⁴

2.8.2 FPU Data Formats

The FPU defines the following data formats:

- 32-bit single-precision floating point (.fmt type *S*)
- 32-bit single-precision floating point paired-single (.fmt type *PS*)⁴
- 64-bit double-precision floating point (.fmt type *D*)
- 32-bit Word fixed point (.fmt type *W*)

4. The CPU Doubleword and FPU floating point paired-single and Long fixed point data formats are available in an implementation that includes a 64-bit floating point unit, whether MIPS32, MIPS64, microMIPS32, or microMIPS64.

- 64-bit Long fixed point (.fmt type *L*)⁴

2.8.3 Coprocessors (CP0-CP3)

The MIPS Architecture defines four coprocessors (designated CP0, CP1, CP2, and CP3):

- Coprocessor 0 (**CP0**) is incorporated on the CPU chip and supports the virtual memory system and exception handling. CP0 is also referred to as the *System Control Coprocessor*.
- Coprocessor 1 (**CP1**) is reserved for the floating point coprocessor, the FPU.
- Coprocessor 2 (**CP2**) is available for specific implementations.
- Coprocessor 3 (**CP3**) is reserved for the floating point unit .

CP0 translates virtual addresses into physical addresses, manages exceptions, and handles switches between kernel, supervisor, and user states. CP0 also controls the cache subsystem, as well as providing diagnostic control and error recovery facilities. The architectural features of CP0 are defined in Volume III.

2.8.4 CPU Registers

The microMIPS64 Architecture defines the following CPU registers:

- 32 64-bit general purpose registers (GPRs)
- a pair of special-purpose registers to hold the results of integer multiply, divide, and multiply-accumulate operations (*HI* and *LO*)
- a special-purpose program counter (*PC*), which is affected only indirectly by certain instructions - it is not an architecturally-visible register.

A MIPS64 processor always produces a 64-bit result, even for those instructions which are architecturally defined to operate on 32 bits. Such instructions typically sign-extend their 32-bit result into 64 bits. In so doing, 32-bit programs work as expected, even though the registers are actually 64 bits wide rather than 32.

2.8.4.1 CPU General-Purpose Registers

Two of the CPU general-purpose registers have assigned functions:

- *r0* is hard-wired to a value of zero, and can be used as the target register for any instruction whose result is to be discarded. *r0* can also be used as a source when a zero value is needed.
- *r31* is the destination register used by JAL, BLTZAL, BLTZALL, BGEZAL, and BGEZALL without being explicitly specified in the instruction word. Otherwise *r31* is used as a normal register.

The remaining registers are available for general-purpose use.

The microMIPS architectures include 16-bit sized instructions. Most of these 16-bit instructions use 3-bit register specifier fields instead of the 5-bit register specifier fields used by most of the 32-bit instructions. Due to these smaller register specifier fields, such instructions can only access 8 of the 32 GPRs. The accessible sets of registers are described in VolumeII-B: The microMIPS Instruction Set. There are also 16-bit move and add instructions which

can directly access all 32 GPRs. In addition, specific instructions implicitly reference r29 (conventionally used as the stack pointer), r28 (conventionally used as the global pointer), and the program counter.

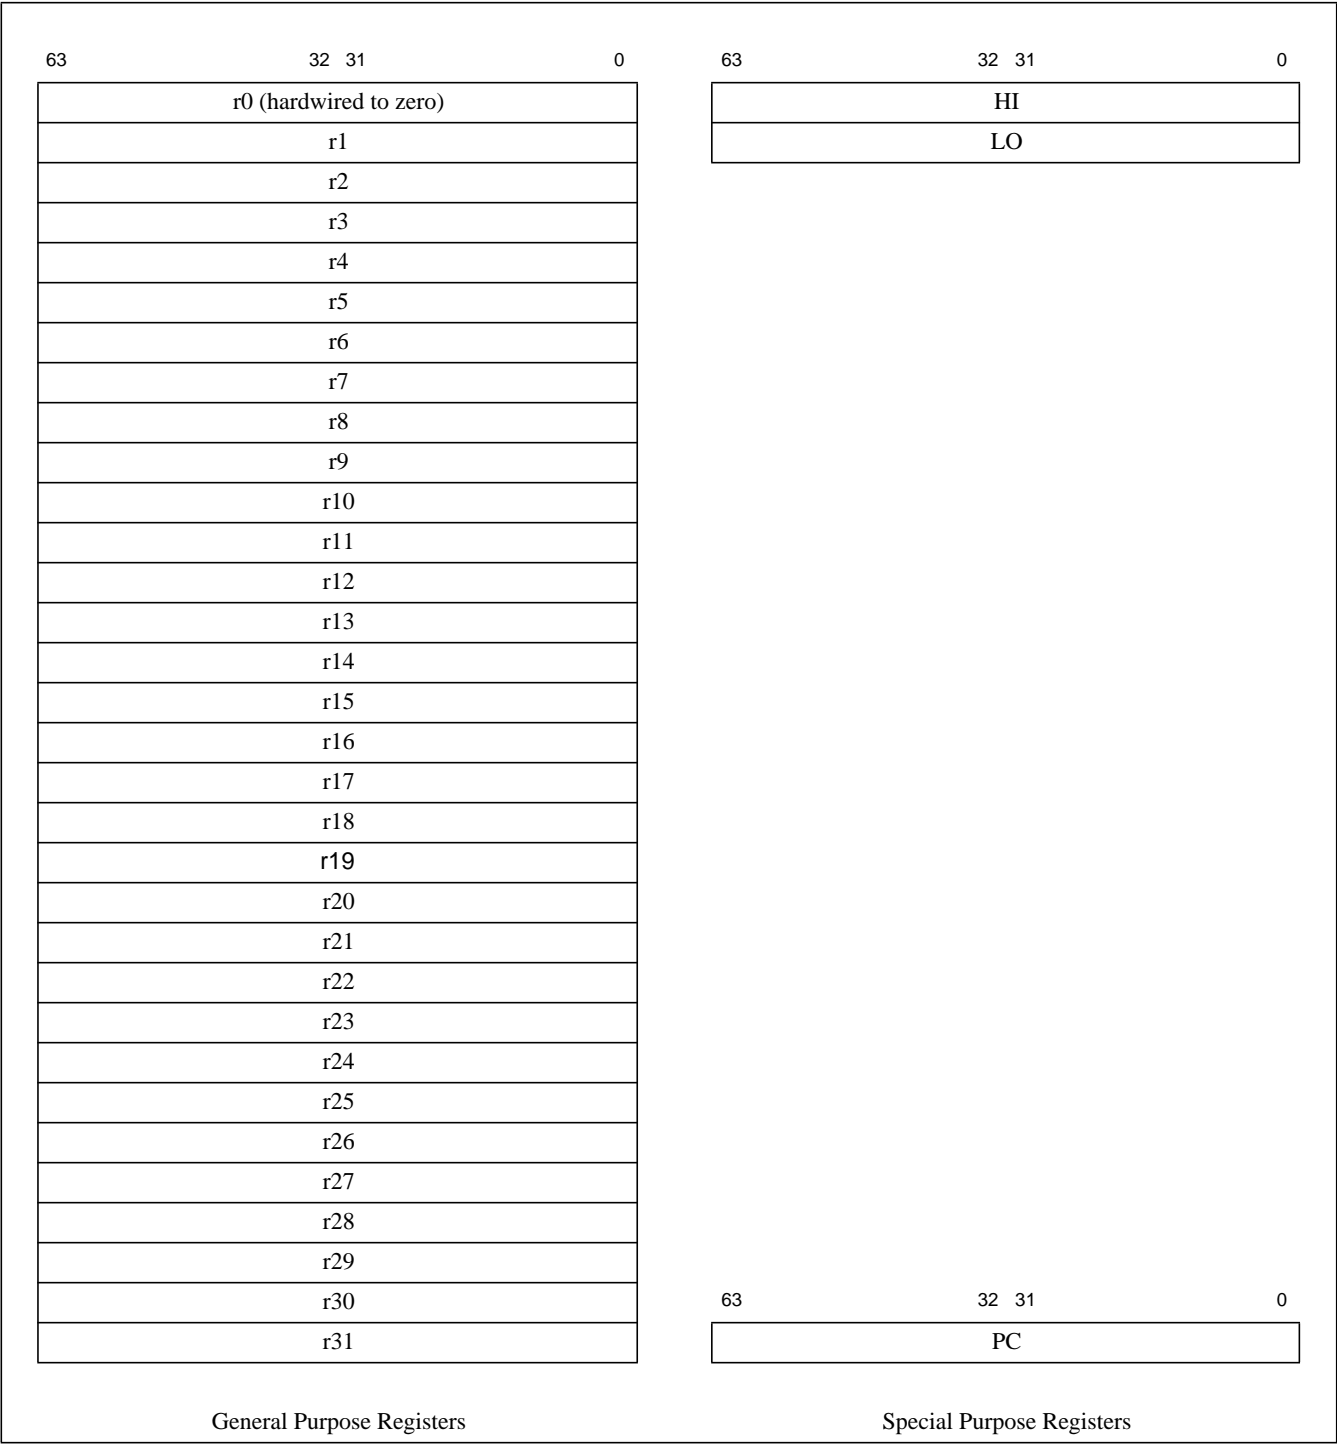
2.8.4.2 CPU Special-Purpose Registers

The CPU contains three special-purpose registers:

- *PC*—Program Counter register
- *HI*—Multiply and Divide register higher result
- *LO*—Multiply and Divide register lower result
 - During a multiply operation, the *HI* and *LO* registers store the product of integer multiply.
 - During a multiply-add or multiply-subtract operation, the *HI* and *LO* registers store the result of the integer multiply-add or multiply-subtract.
 - During a division, the *HI* and *LO* registers store the quotient (in *LO*) and remainder (in *HI*) of integer divide.
 - During a multiply-accumulate, the *HI* and *LO* registers store the accumulated result of the operation.

Figure 2.8 shows the layout of the CPU registers.

Figure 2.8 CPU Registers



2.8.5 FPU Registers

The microMIPS64 Architecture defines the following FPU registers:

- 32 floating point registers (FPRs). These registers are 32 bits wide in a 32-bit FPU and 64 bits wide on a 64-bit FPU.
- Five FPU control registers are used to identify and control the FPU.
- Eight floating point condition codes that are part of the *FCSR* register

A 64-bit floating point unit is optional on implementations of both the microMIPS32 and microMIPS64 Architectures. As of Release 5 of the Architecture, if floating point is implemented then FR=1 is required. I.e. the 64-bit FPU, with the FR=1 64-bit FPU register model, is required. The FR=0 32-bit FPU register model continues to be required.

A 32-bit floating point unit contains 32 32-bit FPRs, each of which is capable of storing a 32-bit data type. Double-precision (type D) data types are stored in even-odd pairs of FPRs, and the long-integer (type L) and paired single (type PS) data types are not supported. [Figure 2.9](#) shows the layout of these registers.

A 64-bit floating point unit contains 32 64-bit FPRs, each of which is capable of storing any data type. For compatibility with 32-bit FPUs, the FR bit in the CP0 *Status* register is used processor that supports a 64-bit FPU to configure the FPU in a mode in which the FPRs are treated as 32 32-bit registers, each of which is capable of storing only 32-bit data types. In this mode, the double-precision floating point (type D) data type is stored in even-odd pairs of FPRs, and the long-integer (type L) and paired single (type PS) data types are not supported.

[Figure 2.10](#) shows the layout of the FPU Registers when the FR bit in the CP0 Status register is 1; [Figure 2.11](#) shows the layout of the FPU Registers when the FR bit in the CP0 Status register is 0.

Figure 2.9 FPU Registers for a 32-bit FPU

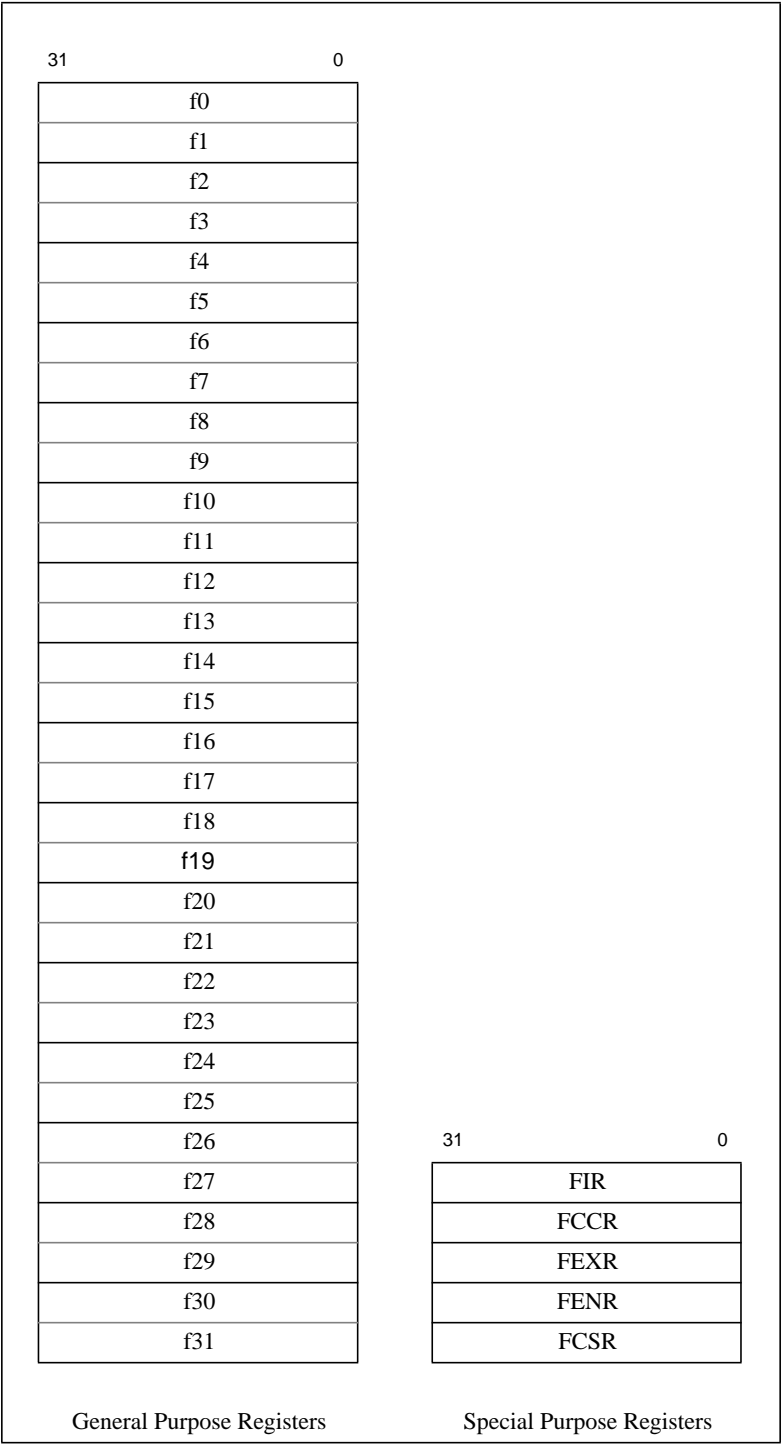


Figure 2.10 FPU Registers for a 64-bit FPU if Status_{FR} is 1

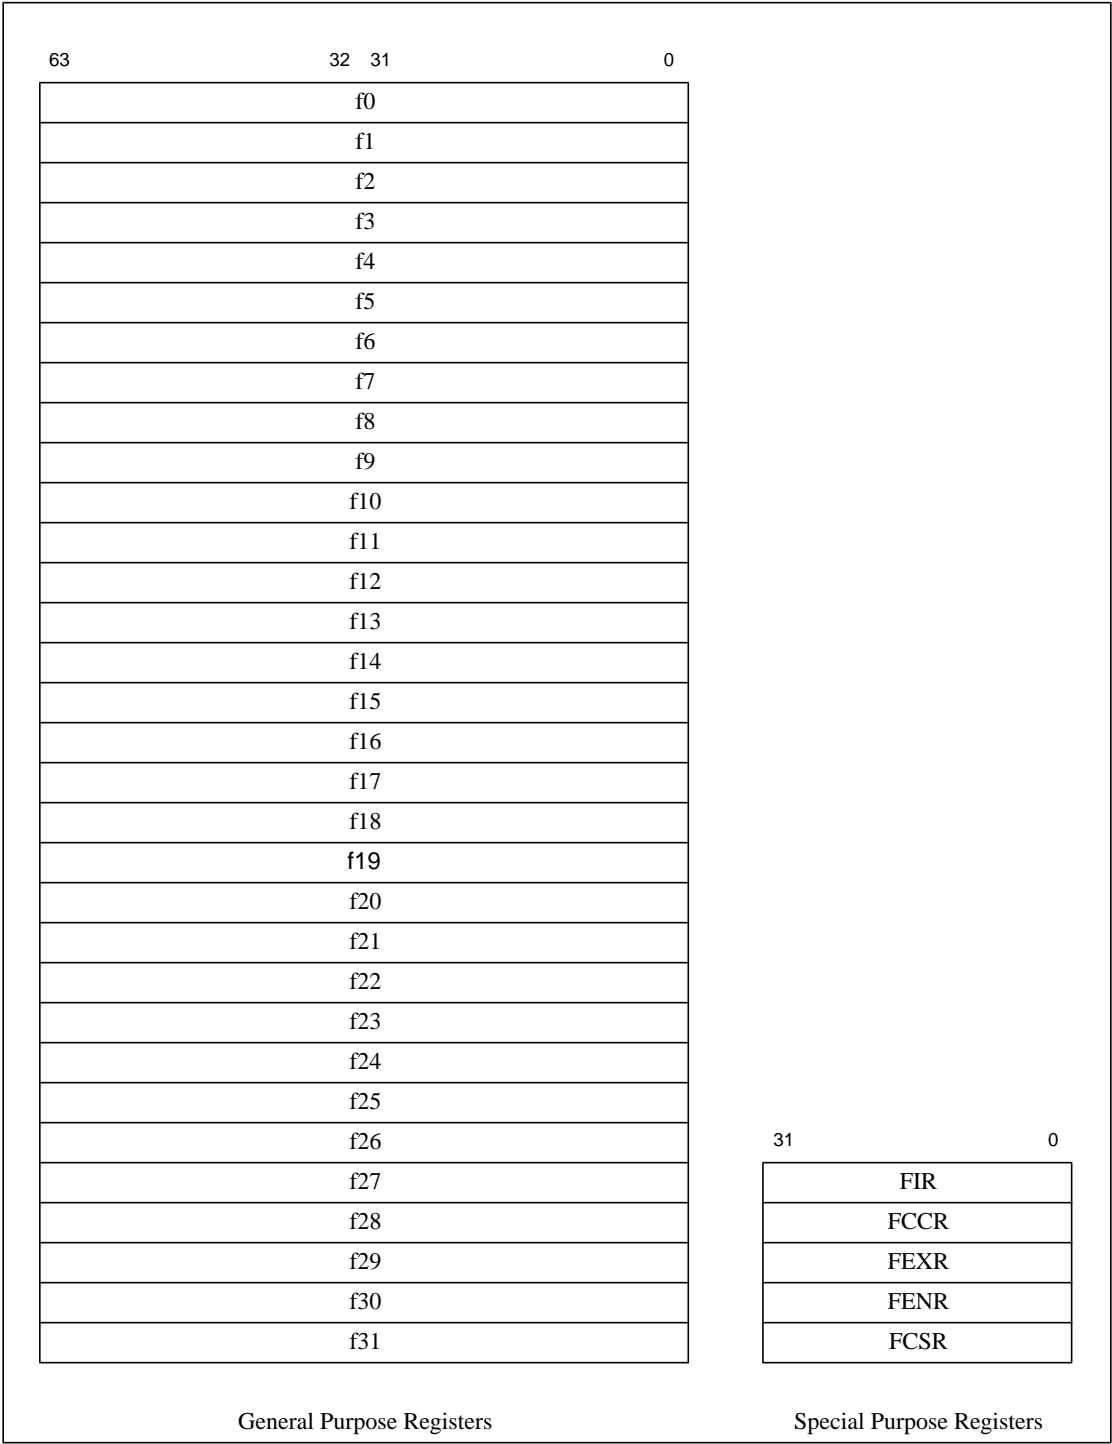
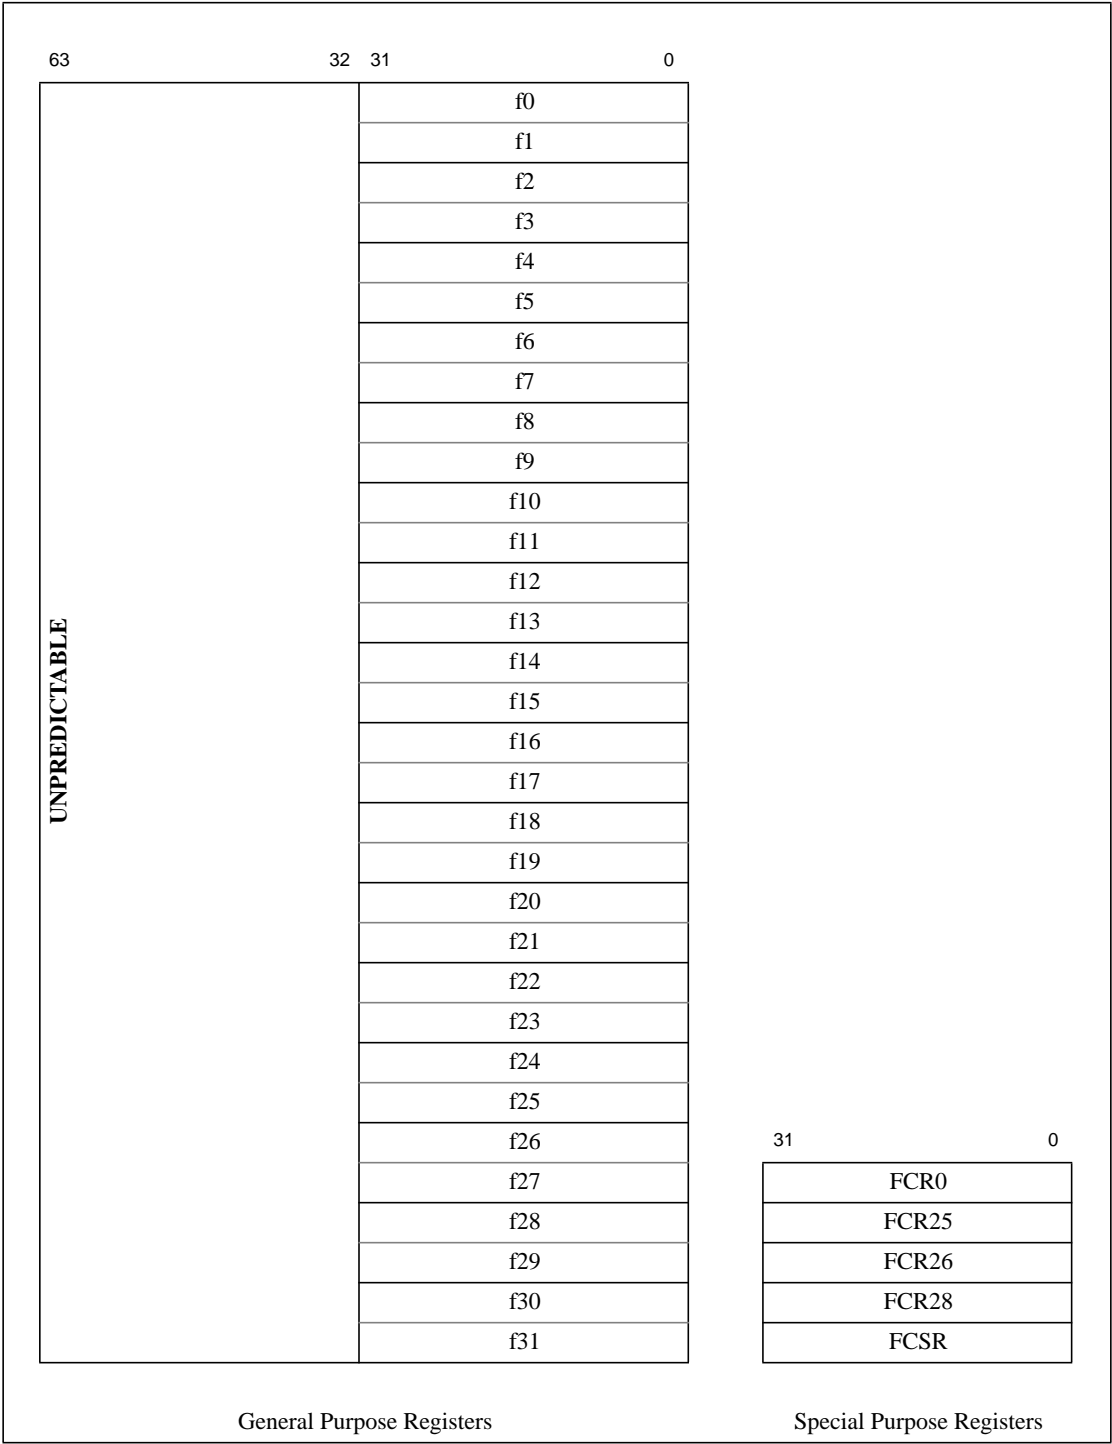


Figure 2.11 FPU Registers for a 64-bit FPU if Status_{FR} is 0



2.8.6 Byte Ordering and Endianness

Bytes within larger CPU data formats—halfword, word, and doubleword—can be configured in either big-endian or little-endian order, as described in the following subsections:


- [Big-Endian Order](#)
- [Little-Endian Order](#)
- [MIPS Bit Endianness](#)

Endianness defines the location of byte 0 within a larger data structure (in this book, bits are always numbered with 0 on the right). Figures 2.12 and 2.13 show the ordering of bytes within words and the ordering of words within multiple-word structures for both big-endian and little-endian configurations.

2.8.6.1 Big-Endian Order

When configured in **big-endian order**, byte 0 is the most-significant (left-hand) byte. [Figure 2.12](#) shows this configuration.


Figure 2.12 Big-Endian Byte Ordering



2.8.6.2 Little-Endian Order

When configured in **little-endian order**, byte 0 is always the least-significant (right-hand) byte. [Figure 2.13](#) shows this configuration.

Figure 2.13 Little-Endian Byte Ordering



2.8.6.3 MIPS Bit Endianness

In this book, bit 0 is always the least-significant (right-hand) bit. Although no instructions explicitly designate bit positions within words, MIPS bit designations are always little-endian.

[2.14](#) shows big-endian and [2.15](#) shows little-endian byte ordering in doublewords.

Figure 2.14 Big-Endian Data in Doubleword Format


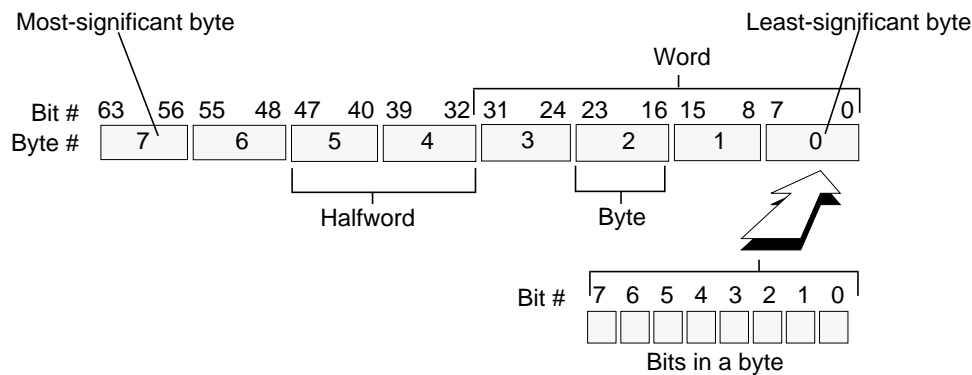


Figure 2.15 Little-Endian Data in Doubleword Format



2.8.7 Memory Alignment

Prior to Release 5 the MIPS architectures required “natural” alignment of memory operands for most memory operations, as described in section 2.8.7.1 “Addressing Alignment Constraints” below. Instructions such as LWL and LWR, described in section 2.8.7.2 “Unaligned Loads and Stores”, were provided so that unaligned accesses could be performed via instruction sequences.

In Release 5 of the Architecture the MSA (MIPS SIMD Architecture) supports 128 bit memory accesses, and does NOT require these to be naturally aligned. The behavior, semantics, and architecture specifications of such mis-aligned accesses are described in an appendix to Volume II of the architecture specification.

2.8.7.1 Addressing Alignment Constraints

The CPU uses byte addressing for halfword, word, and doubleword accesses with the following alignment constraints:

- Halfword accesses must be aligned on an even byte boundary (0, 2, 4...).
- Word accesses must be aligned on a byte boundary divisible by four (0, 4, 8...).
- Doubleword accesses must be aligned on a byte boundary divisible by eight (0, 8, 16...).

2.8.7.2 Unaligned Loads and Stores

The following instructions load and store words that are not aligned on word (W) or doubleword (D) boundaries:

Table 2.1 Unaligned Load and Store Instructions

Alignment	Instructions	Instruction Set
Word	LWL, LWR, SWL, SWR	microMIPS32 ISA
Doubleword	LDL, LDR, SDL, SDR	microMIPS64 ISA

2.16 show a big-endian access of a misaligned word that has byte address 3, and 2.17 shows a little-endian access of a misaligned word that has byte address 1.⁵

Figure 2.16 Big-Endian Misaligned Word Addressing





Figure 2.17 Little-Endian Misaligned Word Addressing



2.8.8 Memory Access Types

MIPS systems provide several *memory access types*. These are characteristic ways to use physical memory and caches to perform a memory access.

The **memory access type** is identified by the Cacheability and Coherency Attribute (CCA) bits in the TLB entry for each mapped virtual page. The access type used for a location is associated with the virtual address, not the physical address or the instruction making the reference. Memory access types are available for both uniprocessor and multi-processor (MP) implementations.

All implementations must provide the following memory access types:

- Uncached

5. These two figures show left-side misalignment.

- Cached

Implementations may include this optional memory access type:

- Uncached Accelerated

These memory access types are described in the following sections:

- [Uncached Memory Access](#)
- [Cached Memory Access](#)
- [Uncached Accelerated Memory Access](#)

2.8.8.1 Uncached Memory Access

In an *uncached* access, physical memory resolves the access. Each reference causes a read or write to physical memory. Caches are neither examined nor modified.

2.8.8.2 Cached Memory Access

In a *cached* access, physical memory and all caches in the system containing a copy of the physical location are used to resolve the access. A copy of a location is coherent if the copy was placed in the cache by a *cached coherent* access; a copy of a location is noncoherent if the copy was placed in the cache by a *cached noncoherent* access. (Coherency is dictated by the system architecture, not the processor implementation.)

Caches containing a coherent copy of the location are examined and/or modified to keep the contents of the location coherent. It is not possible to predict whether caches holding a noncoherent copy of the location will be examined and/or modified during a *cached coherent* access.

Prefetches for data and instructions are allowed. Speculative prefetching of data that may never be used or instructions which may never be executed are allowed.

2.8.8.3 Uncached Accelerated Memory Access

Previous to the 3.5 version of this specification, the behavior of Uncached Accelerated Memory Access type was not architecturally defined, but rather was implementation-specific behavior. In the 3.5 version of this specification, the behavior of the Uncached Accelerated is now architecturally defined. This access type is optional.

In an *uncached accelerated* access, physical memory resolves the access. Each reference causes a read or write to physical memory. Caches are neither examined nor modified.

In *uncached* access, each store instruction causes a separate, unique request to physical memory.

In MIPS CPUs, writes are allowed to be buffered within the CPU. Write buffers are usually of cache-line in size. Usually, if there is sufficient data within the write buffer, the data is sent in one burst transaction for higher efficiency.

In *uncached accelerated* access, the data from multiple store instructions can be sent together to the physical memory in one burst transaction. This is achieved by using write buffers to gather the data from multiple store instructions before sending out the memory request.

Data from store instructions using *uncached accelerated* access are kept in the buffer under these rules:

- Buffering can start on any byte address.
- Data is placed into the buffer obeying full byte addressing.
- Data is placed into the buffer for any request size - byte, half-word, word, double-word and the 3, 5-7 byte sizes allowed by SWR/SWL/SDR/SDL instructions.
- A byte can be over-written with new data before the buffer data is flushed out of the core.
- Multiple buffers (each holding data from multiple store instructions) can be active at one time.

The *uncached accelerated* data within the write-buffer is sent to physical memory under these rules:

- As a consequence of a SYNC instruction being executed. All *uncached accelerated* data within all write buffers is sent to physical memory in this situation.
- If a write-buffer is entirely full with *uncached accelerated* data. Normally, this means an entire cache-line of *uncached accelerated* data is held within the buffer.
- If the target address of any load instruction matches the address of any *uncached accelerated* data within the write buffer.
- If the target address of any store instruction using any other type of access type matches the address of any *uncached accelerated* data within the write buffer.
- As a consequence of a non-coherent SYNCI instruction being executed. All *uncached accelerated* data within all write buffers is sent to physical memory in this situation.
- If the target address of a PREF Nudge operation matches the address of any *uncached accelerated* data within the write buffer.
- All write-buffers capable of holding *uncached accelerated* data are already active and another store instruction using *uncached accelerated* access is executed and whose target address does not match any of these write-buffers. In this case, at least one of the write-buffers must be emptied to physical memory to make space for the new store data.

2.8.9 Implementation-Specific Access Types

An implementation may provide memory access types other than *uncached* or *cached*. Implementation-specific documentation accompanies each processor, and defines the properties of the new access types and their effect on all memory-related operations.

2.8.10 Cacheability and Coherency Attributes and Access Types

Memory access types are specified by architecturally-defined and implementation-specific Cacheability and Coherency Attribute bits (CCAs) kept in TLB entries.

Slightly different cacheability and coherency attributes such as “cached coherent, update on write” and “cached coherent, exclusive on write” can map to the same memory access type; in this case they both map to *cached coherent*. In order to map to the same access type, the fundamental mechanisms of both CCAs must be the same.

When the operation of the instruction is affected, the instructions are described in terms of memory access types. The load and store operations in a processor proceed according to the specific *CCA* of the reference, however, and the pseudocode for load and store common functions uses the *CCA* value rather than the corresponding memory access type.

2.8.11 Mixing Access Types

It is possible to have more than one virtual location mapped to the same physical location (known as **aliasing**). The memory access type used for the virtual mappings may be different, but it is not generally possible to use mappings with different access types at the same time.

For all accesses to virtual locations with the *same* memory access type, a processor executing load and store instructions on a physical location must ensure that the instructions occur in proper program order.

A processor can execute a load or store to a physical location using one access type, but any subsequent load or store to the same location using a different memory access type is **UNPREDICTABLE**, unless a privileged instruction sequence to change the access type is executed between the two accesses. Each implementation has a privileged implementation-specific mechanism to change access types.

The memory access type of a location affects the behavior of I-fetch, load, store, and prefetch operations to that location. In addition, memory access types affect some instruction descriptions. Load Linked (LL, LLD) and Store Conditional (SC, SCD) have defined operation only for locations with *cached* memory access type.

2.8.12 Instruction Fetches

2.8.12.1 Instruction fields layout

For MIPS instructions, the layout of the bit fields within the instructions stays the same regardless of the endianness mode in which the processor is executing. The MIPS architecture only uses Little-Endian bit orderings. Bit 0 of an instruction is always the right-most bit within the instruction while bit 31 is always the left-most bit within a 32-bit instruction. The major opcode is always the left-most 6 bits within the instruction.

2.8.12.2 microMIPS32 and microMIPS64 Instruction placement and endianness

For the microMIPS32 and microMIPS64 architectures, instructions are either 16 or 32 bits. All instructions are aligned to 2-byte boundaries in memory (address bits 0 are 0b0). Instructions of 32-bit size can cross 4-byte boundaries.

Instruction words are always placed in memory according to the endianness.

Figure 2.18 shows an example where the width of external memory is 64-bits (two words) and the processor is executing in little-endian mode and the instructions are placed in memory for little-endian execution. In this case, the less significant address is the the right-most word of the dword while the more significant address is the left-most word within the dword. This example shows a 32-bit instruction crossing a 4-byte (word) boundary.



Figure 2.18 Three instructions placed in a 64-bit wide, little-endian memory

Figure 2.19 shows the equivalent Big-Endian example where the less significant address refers to the left-most word within the dword and the more significant address refers to the right-most word within the dword. In both BE and LE examples, the bit locations within the instruction words has not changed. The location of the major opcode is always at the left-most bits within the word. This example shows a 32-bit instruction which is aligned to a 4-byte (word) boundary.

Figure 2.19 Three instructions placed in a 64-bit wide, big-endian memory

2.8.12.3 Instruction fetches using uncached access to memory without side-effects

Memory regions having no access side-effects can be read an infinite amount of times without changing the value received. For such regions accessed with uncached instruction fetches, the following behaviors are allowed:

It is allowed for the fetch transfer size for uncached memory access to be larger than one instruction word. In this case, it is implementation specific whether multiple instruction fetches are done to the same memory location. It is not required for the processor to have a register to buffer the un-used instructions of the transfer for subsequent execution.

Speculative instruction fetches are allowed. [Table 2.2](#) list some types of speculative instruction fetches.

Table 2.2 Speculative instruction fetches

Sequential instructions located after branch/jump fetched before the branch/jump taken/not-taken decision has been determined.
Predicted branch/jump target addresses fetched before branch/jump taken/not-taken decision has been determined or before when target address has been calculated.
Predicted jump target register values before target register has been read.
Predicted return addresses before return register has been read.
Any other type of prefetching ahead of execution.

2.8.12.4 Instruction fetches using uncached access to memory with side-effects

Access side-effects for a memory region might include FIFO behavior, stack behavior or have location-specific behavior (one memory location defining the behavior of another memory location). For such regions accessed with uncached instruction fetches, these are the architectural requirements:

The transfer size can only be one instruction word per instruction fetch.

Speculative instruction fetches are not allowed. The types of instruction fetches listed in [Table 2.2](#) are not allowed.

The architecture defines this memory segment with access side-effects:

- EJTAG Debug Memory space (dmseg). Please refer to MIPS document - *MD00047 EJTAG Specification*.

Beyond this defined segment, the system programmer/designer is reminded that it is possible to memory map an IO device with access side-effects to any uncached memory location, even within segments which the architecture does not define to have access side-effects. For that reason, any implementation which allows behaviors listed in [2.8.12.3 “Instruction fetches using uncached access to memory without side-effects”](#) should restrict software from executing code within any memory region with side-effects.

2.8.12.5 Instruction fetches using cacheable access to memory

The minimum transfer size for cacheable access is one cacheline. The transfer size may be multiple whole cachelines.

Speculative accesses to cacheable memory spaces are allowed as cacheable memory spaces are defined to have no access side-effects. [Table 2.2](#) list some types of speculative instruction fetches.

2.8.12.6 Instruction fetches and exceptions

Precise exception model for instruction fetches

The MIPS architecture uses the precise exception model for instruction fetches. A precise exception means that for an instruction-sourced exception, the cause of an exception is reported on the exact instruction which the processor has attempted to execute and has caused the exception.

It is not allowed to report an exception for an instruction which could not be executed due to program control flow. For example, if a branch/jump is taken and the instruction after the branch is not to be executed, the address checks (alignment, MMU match/validity, access privilege) for that not-to-be-executed instruction may not generate any exception.

Instruction fetch exceptions on branch delay-slots

For instructions occupying a branch delay-slot, any exceptions, including those generated by the fetch of that instruction, should report the exception results so that the branch can be correctly replayed upon return from the exception handler.

2.8.12.7 Self-Modified Code

When the processor writes memory with new instructions at run-time, there are some software steps that must be taken to ensure that the new instructions are fetched properly.

1. The path of instruction fetches to external memory may not be the same as the path of data loads/stores to external memory (this feature is known as a Harvard architecture). The new instructions must be flushed out to the first level of the memory hierarchy which is shared by both the instruction fetches and the data load/stores.
2. The processor must wait until all of the new instructions have actually been written to this shared level of the memory hierarchy.
3. If there are caches which hold instructions between that first shared level of memory hierarchy and the processor pipeline, any stale instructions within those caches must be first invalidated before executing the new instructions.
4. Some processors might implement some type of instruction prefetching. Precautions must be used to ensure that the prefetching does not interfere with the previous steps.

Modules and Application Specific Extensions

This section gives an overview of the Modules and Architecture Specific Extensions that are supported by the MIPS Architecture Family.

3.1 Description of Optional Components

As the MIPS architecture is adopted into a wider variety of markets, the need to extend this architecture in different directions becomes more and more apparent. Therefore various optional components are provided for use with the base ISAs (MIPS32/MIPS64 and microMIPS32/microMIPS64).

These optional components are licensed to MIPS architecture licensees in two different ways:


1. **Modules** - these optional components are part of the Base Architecture (Revision 5 and newer). If a company has licensed one of the base architectures from MIPS Technologies, then that company has also rights to implement any of the associated modules of that base architecture.
2. **Application Specific Extensions** - these optional components are sold as separate architecture products from MIPS Technologies.

The Modules and ASEs are both optional, so the architecture is not permanently bound to support them and the ASEs are used only as needed.

Extensions to the ISA are driven by the requirements of the computer segment, or by customers whose focus is primarily on performance. A Module or ASE can be used with the appropriate ISA to meet the needs of a specific application or an entire class of applications.

Figure 3.1 shows how ASEs interrelate with the microMIPS32 and microMIPS64 ISAs.

Figure 3.1 microMIPS ISAs and ASEs



microMIPS32 Architecture is a strict subset of the microMIPS64 Architecture. ASEs are applicable to one or both of the base architectures as dictated by market need and the requirements placed on the base architecture by the ASE definition.

3.2 List of Application Specific Instructions

As of the publishing date of this document, the following Application Specific Extensions were supported by the architecture.

Component	Module or ASE	Supported Base Architectures	Use
MIPS-3D®	ASE	microMIPS32 or microMIPS64	Geometry Processing
SmartMIPS®	ASE	microMIPS32	Smart Cards and Smart Objects
MIPS® DSP	Module	microMIPS32 or microMIPS64	Signal Processing
MIPS® MT	Module	microMIPS32 or microMIPS64	Multi-Threading
MCU	ASE	microMIPS32 or microMIPS64	Fast Interrupt Response & I/O register programming
VZE	Module	microMIPS32 or microMIPS64	Hardware Support for Virtualization
MSA	Module	microMIPS32 or microMIPS64	SIMD support

3.2.1 The MDMX™ Application Specific Extension to the microMIPS64 Architectures

The MIPS Digital Media Extension (MDMX) provides video, audio, and graphics pixel processing through vectors of small integers. Although not a part of the MIPS ISA, this extension is included for informational purposes. Volume IV-b of this document set describes the MDMX ASE.

3.2.2 The MIPS-3D® Application Specific Extension to the microMIPS Architecture

The MIPS-3D ASE provides enhanced performance of geometry processing calculations by building on the paired single floating point data type, and adding specific instructions to accelerate computations on these data types. Volume IV-c of this document set describes the MIPS-3D ASE. Because the MIPS-3D ASE requires a 64-bit floating point unit, it is only available with processor that includes a 64-bit FPU.

3.2.3 The SmartMIPS® Application Specific Extension to the microMIPS32 Architecture

The SmartMIPS ASE extends the microMIPS32 Architectures with a set of new and modified instruction designed to improve the performance and reduce the memory consumption of MIPS-based smart card or smart object systems. Because the SmartMIPS ASE requires the MIPS32 Architecture, it is not discussed in this document set.

3.2.4 The MIPS® DSP Module to the MIPS Architecture

The MIPS DSP Module provides enhanced performance of signal-processing applications by providing computational support for fractional data types, SIMD, saturation, and other elements that are commonly used in such applications. Volume IV-e of this document set describes the MIPS DSP Module.

3.2.5 The MIPS® MT Module to the MIPS Architecture

The MIPS MT Module provides the architecture to support multi-threaded implementations of the Architecture. This includes support for both virtual processors and lightweight thread contexts. Volume IV-f of this document set describes the MIPS MT Module.

3.2.6 The MIPS® MCU Application Specific Extension to the MIPS Architecture

The MIPS MCU ASE provides enhanced handling of memory-mapped I/O registers and lower interrupt latencies. Volume IV-g of this document set describes the MIPS MCU ASE.

3.2.7 The MIPS® Virtualization Module to the MIPS Architecture

The MIPS Virtualization Module provides hardware acceleration of virtualization of Operating Systems. Volume IV-i of this document set describes the MIPS VZ Module.

3.2.8 The MIPS® SIMD Architecture Module to the MIPS Architecture

The MIPS SIMD Architecture Module provides high performance parallel processing of vector operations through the use of 128-bit wide vector registers. Volume IV-j of this document set describes the MIPS MSA Module.

Overview of the CPU Instruction Set

This chapter gives an overview of the CPU instructions, including a description of CPU instruction formats. An overview of the FPU instructions is given in [Chapter 5, “Overview of the FPU Instruction Set”](#) on page 68.

4.1 CPU Instructions, Grouped By Function

CPU instructions are organized into the following functional groups:

- Load and store
- Computational
- Jump and branch
- Miscellaneous
- Coprocessor

Each instruction is either 16- bits or 32-bits long.

4.1.1 CPU Load and Store Instructions

MIPS processors use a load/store architecture; all operations are performed on operands held in processor registers and main memory is accessed only through load and store instructions.

4.1.1.1 Types of Loads and Stores

There are several different types of load and store instructions, each designed for a different purpose:

- Transferring variously-sized fields (for example, LB, SW)
- Trading transferred data as signed or unsigned integers (for example, LHU)
- Accessing unaligned fields (for example, LWR, SWL)
- Selecting the addressing mode (for example, SDXC1, in the FPU)
- Atomic memory update (read-modify-write: for instance, LL/SC)

Regardless of the byte ordering (big- or little-endian), the address of a halfword, word, or doubleword is the lowest byte address among the bytes forming the object:

- For big-endian ordering, this is the most-significant byte.

- For a little-endian ordering, this is the least-significant byte.

Refer to “Byte Ordering and Endianness” on page 37 for more information on big-endian and little-endian data ordering.

4.1.1.2 Load and Store Access Types

Tables 4.1 and 4.2 list the data sizes that can be accessed through CPU load and store operations. These tables also indicate the particular ISA within which each operation is defined.

Table 4.1 Load and Store Operations Using Register + Offset Addressing Mode

Data Size	CPU			Coprocessors 1 and 2	
	Load Signed	Load Unsigned	Store	Load	Store
Byte	microMIPS32	microMIPS32	microMIPS32		
Halfword	microMIPS32	microMIPS32	microMIPS32		
Word	microMIPS32	microMIPS64	microMIPS32	microMIPS32	microMIPS32
Doubleword (CPU)	microMIPS64		microMIPS64		
Doubleword (FPU)				microMIPS32	microMIPS32
Unaligned word	microMIPS32		microMIPS32		
Unaligned doubleword	microMIPS64		microMIPS64		
Linked word (atomic modify)	microMIPS32		microMIPS32		
Linked doubleword (atomic modify)	microMIPS64		microMIPS64		

Table 4.2 FPU Load and Store Operations Using Register + Register Addressing Mode

Floating Point Coprocessor Only		
Data Size	Load	Store
Word	microMIPS32	microMIPS32
Doubleword	microMIPS32	microMIPS32
Unaligned Doubleword Indexed	microMIPS32	microMIPS32

4.1.1.3 List of CPU Load and Store Instructions

The following data sizes (as defined in the *AccessLength* field) are transferred by CPU load and store instructions:

- Byte
- Halfword
- Word
- Doubleword

Signed and unsigned integers of different sizes are supported by loads that either sign-extend or zero-extend the data loaded into the register.

Table 4.3 lists aligned CPU load and store instructions, while unaligned loads and stores are listed in Table 4.4. Each table also lists the MIPS ISA within which an instruction is defined.

Table 4.3 Aligned CPU Load/Store Instructions

Mnemonic	Instruction	Defined in MIPS ISA
LB	Load Byte	microMIPS32
LBE	Load Byte EVA	microMIPS32
LBU	Load Byte Unsigned	microMIPS32
LBUE	Load Byte Unsigned EVA	microMIPS32
LBU16	Load Byte Unsigned (16-bit Instruction Size)	microMIPS32
LD	Load Doubleword	microMIPS64
LDM	Load Double Multiple	microMIPS64
LDP	Load Double Pair	microMIPS32
LH	Load Halfword	microMIPS32
LHE	Load Halfword EVA	microMIPS32
LHU	Load Halfword Unsigned	microMIPS32
LHUE	Load Halfword Unsigned EVA	microMIPS32
LHU16	Load Halfword Unsigned (16-bit Instruction Size)	microMIPS32
LW	Load Word	microMIPS32
LWE	Load Word EVA	microMIPS32
LW16	Load Word (16-bit Instruction Size)	microMIPS32
LWM32	Load Word Multiple (32-bit Instruction Size)	microMIPS32
LWM16	Load Word Multiple (16-bit Instruction Size)	microMIPS32
LWP	Load Word Pair	microMIPS32
LWGP	Load Word Global Pointer (16-bit Instruction Size)	microMIPS32
LWSP	Load Word Stack Pointer (16-bit Instruction Size)	microMIPS32
LWU	Load Word Unsigned	microMIPS64
SB	Store Byte	microMIPS32
SBE	Store Byte EVA	microMIPS32
SB16	Store Byte (16-bit Instruction Size)	microMIPS32
SD	Store Doubleword	microMIPS64
SDM	Store Double Multiple	microMIPS32
SDP	Store Double Pair	microMIPS32
SH	Store Halfword	microMIPS32
SHE	Store Halfword EVA	microMIPS32
SH16	Store Half (16-bit Instruction Size)	microMIPS32
SW	Store Word	microMIPS32
SWE	Store Word EVA	microMIPS32

Table 4.3 Aligned CPU Load/Store Instructions (Continued)

Mnemonic	Instruction	Defined in MIPS ISA
SW16	Store Word (16-bit Instruction Size)	microMIPS32
SWSP	Store Word Stack Pointer (16-bit Instruction Size)	microMIPS32
SWM32	Store Word Multiple (32-bit Instruction Size)	microMIPS32
SWM16	Store Word Multiple (16-bit Instruction Size)	microMIPS32
SWP	Store Word Pair	microMIPS32

Unaligned words and doublewords can be loaded or stored in just two instructions by using a pair of the special instructions listed in [Table 4.4](#). The load instructions read the left-side or right-side bytes (left or right side of register) from an aligned word and merge them into the correct bytes of the destination register.

Unaligned CPU load and store instructions are listed in [Table 4.4](#), along with the MIPS ISA within which an instruction is defined.

Table 4.4 Unaligned CPU Load and Store Instructions

Mnemonic	Instruction	Defined in MIPS ISA
LDL	Load Doubleword Left	microMIPS64
LDR	Load Doubleword Right	microMIPS64
LWL	Load Word Left	microMIPS32
LWR	Load Word Right	microMIPS32
SDL	Store Doubleword Left	microMIPS64
SDR	Store Doubleword Right	microMIPS64
SWL	Store Word Left	microMIPS32
SWR	Store Word Right	microMIPS32

4.1.1.4 Loads and Stores Used for Atomic Updates

The paired instructions, Load Linked and Store Conditional, can be used to perform an atomic read-modify-write of word or doubleword cached memory locations. These instructions are used in carefully coded sequences to provide one of several synchronization primitives, including test-and-set, bit-level locks, semaphores, and sequencers and event counts. [Table 4.5](#) lists the LL and SC instructions, along with the MIPS ISA within which an instruction is defined.

Table 4.5 Atomic Update CPU Load and Store Instructions

Mnemonic	Instruction	Defined in MIPS ISA
LL	Load Linked Word	microMIPS32
LLD	Load Linked Doubleword	microMIPS64
SC	Store Conditional Word	microMIPS32
SCD	Store Conditional Doubleword	microMIPS64

4.1.1.5 CPU Loads and Stores Using Register + Register Addressing

Figure 4.6 lists the CPU Load/Store instructions which use register + register addressing.

Table 4.6 CPU Load and Store Instructions Using Register + Register Addressing

Mnemonic	Instruction	Defined in MIPS ISA
LWXS	Load Word Indexed-Scaled	microMIPS32

4.1.1.6 Coprocessor Loads and Stores

If a particular coprocessor is not enabled, loads and stores to that processor cannot execute and the attempted load or store causes a Coprocessor Unusable exception. Enabling a coprocessor is a privileged operation provided by the System Control Coprocessor, CP0.

Table 4.7 lists the coprocessor load and store instructions.

Table 4.7 Coprocessor Load and Store Instructions

Mnemonic	Instruction	Defined in MIPS ISA
LDCz	Load Doubleword to Coprocessor-z, z = 1 or 2	microMIPS32
LWCz	Load Word to Coprocessor-z, z = 1 or 2	microMIPS32
SDCz	Store Doubleword from Coprocessor-z, z = 1 or 2	microMIPS32
SWCz	Store Word from Coprocessor-z, z = 1 or 2	microMIPS32

Table 4.8 lists the specific FPU load and store instructions;¹ it also lists the MIPS ISA within which an instruction was first defined.

Table 4.8 FPU Load and Store Instructions Using Register + Register Addressing

Mnemonic	Instruction	Defined in MIPS ISA
LWXC1	Load Word Indexed to Floating Point	microMIPS32
SWXC1	Store Word Indexed from Floating Point	microMIPS32
LDXC1	Load Doubleword Indexed to Floating Point	microMIPS32
SDXC1	Store Doubleword Indexed from Floating Point	microMIPS32
LUXC1	Load Doubleword Indexed Unaligned to Floating Point	microMIPS32
SUXC1	Store Doubleword Indexed Unaligned from Floating Point	microMIPS32

4.1.2 Computational Instructions

This section describes the following:

- [ALU Immediate and Three-Operand Instructions](#)
- [ALU Two-Operand Instructions](#)
- [Shift Instructions](#)

1. FPU loads and stores are listed here with the other coprocessor loads and stores for convenience.

- [Multiply and Divide Instructions](#)

2's complement arithmetic is performed on integers represented in 2's complement notation. These are signed versions of the following operations:

- Add
- Subtract
- Multiply
- Divide

The add and subtract operations labelled “unsigned” are actually modulo arithmetic without overflow detection.

There are also unsigned versions of *multiply* and *divide*, as well as a full complement of *shift* and *logical* operations. Logical operations are not sensitive to the width of the register.

microMIPS32 provided 32-bit integers and 32-bit arithmetic. MIPS64 adds 64-bit integers and provides separate arithmetic and shift instructions for 64-bit operands.

4.1.2.1 ALU Immediate and Three-Operand Instructions

[Table 4.9](#) lists those arithmetic and logical instructions that operate on one operand from a register and the other from a 16-bit *immediate* value supplied by the instruction word. This table also lists the MIPS ISA within which an instruction is defined.

The *immediate* operand is treated as a signed value for the arithmetic and compare instructions, and treated as a logical value (zero-extended to register length) for the logical instructions.

Table 4.9 ALU Instructions With a 16-bit Immediate Operand

Mnemonic	Instruction	Defined in MIPS ISA
ADDI	Add Immediate Word	microMIPS32
ADDIU ¹	Add Immediate Unsigned Word	microMIPS32
ANDI	And Immediate	microMIPS32
DADDI	Doubleword Add Immediate	microMIPS64
DADDIU ¹	Doubleword Add Immediate Unsigned	microMIPS64
LUI	Load Upper Immediate	microMIPS32
ORI	Or Immediate	microMIPS32
SLTI	Set on Less Than Immediate	microMIPS32
SLTIU	Set on Less Than Immediate Unsigned	microMIPS32
XORI	Exclusive Or Immediate	microMIPS32

1. The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow.

Table 4.10 describes those arithmetic and logical instructions with an *immediate* value whose size is not 16-bits.

Table 4.10 Other ALU Instructions With a Immediate Operand

Mnemonic	Instruction	Defined in MIPS ISA
ADDIUPC	Add Immediate Unsigned Word PC-relative	microMIPS32
ADDIUR1SP	Add Immediate Unsigned Word Register with Stack Pointer (16-bit Instruction Size)	microMIPS32
ADDIUR2	Add Immediate Unsigned Word Two Registers (16-bit Instruction Size)	microMIPS32
ADDIUSP	Add Immediate Unsigned to Stack Pointer (16-bit Instruction Size)	microMIPS32
ADDIUS5	Add Immediate Unsigned Word 5-bit Register Select (16-bit Instruction Size)	microMIPS32
ANDI16	Add Unsigned (16-bit Instruction Size)	microMIPS32
LI16	Load Immediate (16-bit Instruction Size)	microMIPS32

Table 4.11 describes ALU instructions that use three operands, along with the MIPS ISA within which an instruction is defined.

Table 4.11 Three-Operand ALU Instructions

Mnemonic	Instruction	Defined in MIPS ISA
ADD	Add Word	microMIPS32
ADDU ¹	Add Unsigned Word	microMIPS32
ADDU16	Add Unsigned Word (16-bit Instruction Size)	microMIPS32
AND	And	microMIPS32
DADD	Doubleword Add	microMIPS64
DADDU ¹	Doubleword Add Unsigned	microMIPS64
DSUB	Doubleword Subtract	microMIPS64
DSUBU ¹	Doubleword Subtract Unsigned	microMIPS64
MOVEP	Move Pair of Registers (16-bit Instruction Size)	microMIPS32
NOR	Nor	microMIPS32
OR	Or	microMIPS32
SLT	Set on Less Than	microMIPS32
SLTU	Set on Less Than Unsigned	microMIPS32
SUB	Subtract Word	microMIPS32
SUBU ¹	Subtract Unsigned Word	microMIPS32
SUBU16	Subtract Unsigned Word (16-bit Instruction Size)	microMIPS
XOR	Exclusive Or	microMIPS32

1. The term “unsigned” in the instruction name is a misnomer; this operation is 32-bit modulo arithmetic that does not trap on overflow.

4.1.2.2 ALU Two-Operand Instructions

Table 4.11 describes ALU instructions that use two operands, along with the MIPS ISA within which an instruction is defined.

Table 4.12 Two-Operand ALU Instructions

Mnemonic	Instruction	Defined in MIPS ISA
AND16	AND (16-bit Instruction Size)	microMIPS32
CLO	Count Leading Ones in Word	microMIPS32
CLZ	Count Leading Zeros in Word	microMIPS32
DCLO	Count Leading Ones in Doubleword	microMIPS64
DCLZ	Count Leading Zeros in Doubleword	microMIPS64
MOVE16	MOVE Register (16-bit Instruction Size)	microMIPS32
NOT16	Invert (16-bit Instruction Size)	microMIPS32
OR16	OR (16-bit Instruction Size)	microMIPS32
XOR16	XOR (16-bit Instruction Size)	microMIPS32

4.1.2.3 Shift Instructions

The ISA defines three types of shift instructions:

- Those that take a fixed shift amount from a 5-bit field in the instruction word (for instance, SLL, SRL)
- Those that take a shift amount from the low-order bits of a general register (for instance, SRAV, SRLV)
- Those that take as shift amount from an encoded 3-bit field in the instruction word (SLL16, SRL16)

The instructions with a fixed shift amount are limited to a 5-bit shift count, so there are separate instructions for doubleword shifts of 0-31 bits (for instance, DSLL) and 32-63 bits (for instance, DSLL32).

Shift instructions are listed in Table 4.13, along with the MIPS ISA within which an instruction is defined.

Table 4.13 Shift Instructions

Mnemonic	Instruction	Defined in MIPS ISA
DROTR	Doubleword Rotate Right	microMIPS64
DROTR32	Doubleword Rotate Right Plus 32	microMIPS64
DROTRV	Doubleword Rotate Right Variable	microMIPS64
DSLL	Doubleword Shift Left Logical	microMIPS64
DSLL32	Doubleword Shift Left Logical + 32	microMIPS64
DSLLV	Doubleword Shift Left Logical Variable	microMIPS64
DSRA	Doubleword Shift Right Arithmetic	microMIPS64
DSRA32	Doubleword Shift Right Arithmetic + 32	microMIPS64
DSRAV	Doubleword Shift Right Arithmetic Variable	microMIPS64
DSRL	Doubleword Shift Right Logical	microMIPS64
DSRL32	Doubleword Shift Right Logical + 32	microMIPS64

Table 4.13 Shift Instructions (Continued)

Mnemonic	Instruction	Defined in MIPS ISA
DSRLV	Doubleword Shift Right Logical Variable	microMIPS64
ROTR	Rotate Word Right	microMIPS32
ROTRV	Rotate Word Right Variable	microMIPS32
SLL	Shift Word Left Logical	microMIPS32
SLL16	Shift Word Left Logical (16-bit Instruction Size)	microMIPS32
SLLV	Shift Word Left Logical Variable	microMIPS32
SRA	Shift Word Right Arithmetic	microMIPS32
SRAV	Shift Word Right Arithmetic Variable	microMIPS32
SRL	Shift Word Right Logical	microMIPS32
SRL16	Shift Word Right Logical (16-bit Instruction Size)	microMIPS32
SRLV	Shift Word Right Logical Variable	microMIPS32

4.1.2.4 Multiply and Divide Instructions

The multiply and divide instructions produce twice as many result bits as is typical with other processors. With one exception, they deliver their results into the *HI* and *LO* special registers. The *MUL* instruction delivers the lower half of the result directly to a GPR.

- **Multiply** produces a full-width product twice the width of the input operands; the low half is loaded into *LO* and the high half is loaded into *HI*.
- **Multiply-Add** and **Multiply-Subtract** produce a full-width product twice the width of the input operations and adds or subtracts the product from the concatenated value of *HI* and *LO*. The low half of the addition is loaded into *LO* and the high half is loaded into *HI*.
- **Divide** produces a quotient that is loaded into *LO* and a remainder that is loaded into *HI*.

The results are accessed by instructions that transfer data between *HI/LO* and the general registers.

Table 4.14 lists the multiply, divide, and *HI/LO* move instructions, along with the MIPS ISA within which an instruction is defined.

Table 4.14 Multiply/Divide Instructions

Mnemonic	Instruction	Defined in MIPS ISA
DDIV	Doubleword Divide	microMIPS64
DDIVU	Doubleword Divide Unsigned	microMIPS64
DIV	Divide Word	microMIPS32
DIVU	Divide Unsigned Word	microMIPS32
DMULT	Doubleword Multiply	microMIPS64
DMULTU	Doubleword Multiply Unsigned	microMIPS64
MADD	Multiply and Add Word	microMIPS32
MADDU	Multiply and Add Word Unsigned	microMIPS32
MFHI	Move From HI	microMIPS32

Table 4.14 Multiply/Divide Instructions (Continued)

Mnemonic	Instruction	Defined in MIPS ISA
MFHI16	Move From HI (16-bit Instruction Size)	microMIPS32
MFLO	Move From LO	microMIPS32
MFLO16	Move From LO(16-bit Instruction Size)	microMIPS32
MSUB	Multiply and Subtract Word	microMIPS32
MSUBU	Multiply and Subtract Word Unsigned	microMIPS32
MTHI	Move To HI	microMIPS32
MTLO	Move To LO	microMIPS32
MUL	Multiply Word to Register	microMIPS32
MULT	Multiply Word	microMIPS32
MULTU	Multiply Unsigned Word	microMIPS32

4.1.3 Jump and Branch Instructions

This section describes the following:

- [Types of Jump and Branch Instructions Defined by the ISA](#)
- [Branch Delays and the Branch Delay Slot](#)
- [Delay Slot Behavior](#)
- [List of Jump and Branch Instructions](#)

4.1.3.1 Types of Jump and Branch Instructions Defined by the ISA

The architecture defines the following jump and branch instructions:

- PC-relative conditional branch
- PC-region unconditional jump
- Absolute (register) unconditional jump
- A set of procedure calls that record a return link address in a general register.

4.1.3.2 Branch Delays and the Branch Delay Slot

All branches have an architectural delay of one instruction. The instruction immediately following a branch is said to be in the **branch delay slot**. If a branch or jump instruction is placed in the branch delay slot, the operation of both instructions is **UNPREDICTABLE**.

By convention, if an exception or interrupt prevents the completion of an instruction in the branch delay slot, the instruction stream is continued by re-executing the branch instruction. To permit this, branches must be restartable; procedure calls may not use the register in which the return link is stored (usually GPR 31) to determine the branch target address.

4.1.3.3 Delay Slot Behavior

There are three versions of branches and jumps; they differ in the manner in which they handle the instruction in the delay slot when the branch is not taken and execution falls through.

- **Branch** and **Jump** instructions execute the instruction in the delay slot.
- **Branch likely** instructions do not execute the instruction in the delay slot if the branch is not taken (they are said to *nullify* the instruction in the delay slot). The microMIPS architectures do not include any branch likely instructions. They are listed here only for historical background.
- **Compact Branches** and **Compact Jumps** which do not execute the instruction in the delay slot. These branches are available in the microMIPS architectures.

4.1.3.4 List of Jump and Branch Instructions

Table 4.15 lists instructions that jump to a procedure call within the current 256 MB-aligned region.

Table 4.16 lists instructions that jump to an absolute address held in a register.

Table 4.15 lists the unconditional jump instructions within a given 256 MByte region. Table 4.17 lists branch instructions that compare two registers before conditionally executing a PC-relative branch. Table 4.18 lists branch instructions that test a register—compare with zero—before conditionally executing a PC-relative branch.

Each table also lists the MIPS ISA within which an instruction is defined.

Table 4.15 Unconditional Jump Within a 256 Megabyte Region

Mnemonic	Instruction	Defined in MIPS ISA
J	Jump	microMIPS32
JAL	Jump and Link	microMIPS32
JALS	Jump and Link, Short Delay-Slot	microMIPS32
JALX	Jump and Link Exchange	microMIPS32

Table 4.16 Unconditional Jump using Absolute Address

Mnemonic	Instruction	Defined in MIPS ISA
JALR	Jump and Link Register	microMIPS32
JALRS	Jump and Link Register, Short Delay-Slot	microMIPS32
JALR16	Jump and Link Register (16-bit Instruction Size)	microMIPS32
JALR16S	Jump and Link Register, Short Delay-Slot (16-bit Instruction Size)	microMIPS32
JALR.HB	Jump and Link Register with Hazard Barrier	microMIPS32
JALRS.HB	Jump and Link Register with Hazard Barrier, Short Delay-Slot	microMIPS32
JR	Jump Register	microMIPS32
JR16	Jump Register (16-bit Instruction Size)	microMIPS32

Table 4.16 Unconditional Jump using Absolute Address

Mnemonic	Instruction	Defined in MIPS ISA
JRADDIUSP	Jump Register, Adjust Stack Pointer (16-bit Instruction Size)	microMIPS32
JRC	Jump Register, Compact (16-bit Instruction Size)	microMIPS32
JR.HB	Jump Register with Hazard Barrier	microMIPS32

Table 4.17 PC-Relative Conditional Branch Instructions Comparing Two Registers

Mnemonic	Instruction	Defined in MIPS ISA
BEQ	Branch on Equal	microMIPS32
BNE	Branch on Not Equal	microMIPS32

Table 4.18 PC-Relative Conditional Branch Instructions Comparing With Zero

Mnemonic	Instruction	Defined in MIPS ISA
BEQZ16	Branch on Equal to Zero (16-bit Instruction Size)	microMIPS32
BEQZC	Branch of Equal to Zero, Compact (16-bit Instruction Size)	microMIPS32
BGEZ	Branch on Greater Than or Equal to Zero	microMIPS32
BGEZAL	Branch on Greater Than or Equal to Zero and Link	microMIPS32
BGEZALS	Branch on Greater Than or Equal to Zero and Link, Short Delay-Slot	microMIPS32
BGTZ	Branch on Greater Than Zero	microMIPS32
BLEZ	Branch on Less Than or Equal to Zero	microMIPS32
BLTZ	Branch on Less Than Zero	microMIPS32
BLTZAL	Branch on Less Than Zero and Link	microMIPS32
BLTZALS	Branch on Less Than Zero and Link, Short Delay-Slot	microMIPS32
BNEZ16	Branch on Not Equal to Zero (16-bit Instruction Size)	microMIPS32
BNEZC	Branch of Not Equal to Zero, Compact (16-bit Instruction Size)	microMIPS32

Table 4.19 PC-relative Unconditional Branch

Mnemonic	Instruction	Defined in MIPS ISA
B16	Unconditional Branch (16-bit Instruction Size)	microMIPS32

4.1.4 Miscellaneous Instructions

Miscellaneous instructions include:

- [Instruction Serialization \(SYNC and SYNCI\)](#)
- [Exception Instructions](#)
- [Conditional Move Instructions](#)
- [Prefetch Instructions](#)

- [NOP Instructions](#)

4.1.4.1 Instruction Serialization (SYNC and SYNCI)

In normal operation, the order in which load and store memory accesses appear to a viewer *outside* the executing processor (for instance, in a multiprocessor system) is not specified by the architecture.

The SYNC instruction can be used to create a point in the executing instruction stream at which the relative order of some loads and stores can be determined: loads and stores executed before the SYNC are completed before loads and stores after the SYNC can start.

The SYNCI instruction synchronizes the processor caches with previous writes or other modifications to the instruction stream.

[Table 4.20](#) lists the synchronization instructions, along with the MIPS ISA within which it is defined.

Table 4.20 Serialization Instruction

Mnemonic	Instruction	Defined in MIPS ISA
SYNC	Synchronize Shared Memory	microMIPS32
SYNCI	Synchronize Caches to Make Instruction Writes Effective	microMIPS32

4.1.4.2 Exception Instructions

Exception instructions transfer control to a software exception handler in the kernel. There are two types of exceptions, *conditional* and *unconditional*. These are caused by the following instructions:

Trap instructions, which cause conditional exceptions based upon the result of a comparison

System call and breakpoint instructions, which cause unconditional exceptions

[Table 4.21](#) lists the system call and breakpoint instructions. [Table 4.22](#) lists the trap instructions that compare two registers. [Table 4.23](#) lists trap instructions, which compare a register value with an *immediate* value.

Each table also lists the MIPS ISA within which an instruction is defined.

Table 4.21 System Call and Breakpoint Instructions

Mnemonic	Instruction	Defined in MIPS ISA
BREAK	Breakpoint	microMIPS32
BREAK16	Breakpoint (16-bit Instruction Size)	microMIPS32
SYSCALL	System Call	microMIPS32

Table 4.22 Trap-on-Condition Instructions Comparing Two Registers

Mnemonic	Instruction	Defined in MIPS ISA
TEQ	Trap if Equal	microMIPS32
TGE	Trap if Greater Than or Equal	microMIPS32
TGEU	Trap if Greater Than or Equal Unsigned	microMIPS32

Table 4.22 Trap-on-Condition Instructions Comparing Two Registers

Mnemonic	Instruction	Defined in MIPS ISA
TLT	Trap if Less Than	microMIPS32
TLTU	Trap if Less Than Unsigned	microMIPS32
TNE	Trap if Not Equal	microMIPS32

Table 4.23 Trap-on-Condition Instructions Comparing an Immediate Value

Mnemonic	Instruction	Defined in MIPS ISA
TEQI	Trap if Equal Immediate	microMIPS32
TGEI	Trap if Greater Than or Equal Immediate	microMIPS32
TGEIU	Trap if Greater Than or Equal Immediate Unsigned	microMIPS32
TLTI	Trap if Less Than Immediate	microMIPS32
TLTIU	Trap if Less Than Immediate Unsigned	microMIPS32
TNEI	Trap if Not Equal Immediate	microMIPS32

4.1.4.3 Conditional Move Instructions

MIPS32 includes instructions to conditionally move one CPU general register to another, based on the value in a third general register. For floating point conditional moves, refer to Chapter 4.

Table 4.24 lists conditional move instructions, along with the MIPS ISA within which an instruction is defined.

Table 4.24 CPU Conditional Move Instructions

Mnemonic	Instruction	Defined in MIPS ISA
MOVF	Move Conditional on Floating Point False	microMIPS32
MOVN	Move Conditional on Not Zero	microMIPS32
MOVT	Move Conditional on Floating Point True	microMIPS32
MOVZ	Move Conditional on Zero	microMIPS32

4.1.4.4 Prefetch Instructions

There are two prefetch advisory instructions:

- One with register+offset addressing (PREF)
- One with register+register addressing (PREFX)

These instructions advise that memory is likely to be used in a particular way in the near future and should be prefetched into the cache. The PREFX instruction is encoded in the FPU *opcode* space, along with the other operations using register+register addressing

Table 4.25 Prefetch Instructions

Mnemonic	Instruction	Addressing Mode	Defined in MIPS ISA
PREF	Prefetch	Register+Offset	microMIPS32
PREFX	Prefetch Indexed	Register+Register	microMIPS32

4.1.4.5 NOP Instructions

The NOP instruction is actually encoded as an all-zero instruction. MIPS processors special-case this encoding as performing no operation, and optimize execution of the instruction. In addition, SSNOP instruction, takes up one issue cycle on any processor, including super-scalar implementations of the architecture.

Table 4.26 lists conditional move instructions, along with the MIPS ISA within which an instruction is defined.

Table 4.26 NOP Instructions

Mnemonic	Instruction	Defined in MIPS ISA
NOP ¹	No Operation	microMIPS32
SSNOP	Superscalar Inhibit NOP	microMIPS32

1. A NOP of instruction size 16-bits is realized through the macro of “MOVE16 \$0,\$0”

4.1.5 Coprocessor Instructions

This section contains information about the following:

- [What Coprocessors Do](#)
- [System Control Coprocessor 0 \(CP0\)](#)
- [Floating Point Coprocessor 1 \(CP1\)](#)
- [Coprocessor Load and Store Instructions](#)

4.1.5.1 What Coprocessors Do

Coprocessors are alternate execution units, with register files separate from the CPU. In abstraction, the MIPS architecture provides for up to four coprocessor units, numbered 0 to 3. Each level of the ISA defines a number of these coprocessors, as listed in Table 4.27.

Table 4.27 Coprocessor Definition and Use in the MIPS Architecture

Coprocessor	microMIPS32	microMIPS64
CP0	Sys Control	Sys Control
CP1	FPU	FPU
CP2	implementation specific	
CP3	See Footnote ¹ MIPS32r1: imp. spec. MIPS32r2: FPU (COP1X)	FPU (COP1X)

1. In Release 1 of the MIPS32 Architecture, Coprocessor 3 was an implementation-specific coprocessor. In the MIPS64 Architecture, and in Release 2 of the MIPS32 Architecture (and subsequent releases) Architectures, it is used exclusively for the floating point unit and is not available for implementation-specific use. Release 1 MIPS32 implementations are encouraged not to use Coprocessor 3 as an implementation-specific coprocessor.

Coprocessor 0 is always used for system control and coprocessor 1 and 3 are used for the floating point unit. Coprocessor 2 is reserved for implementation-specific use.

A coprocessor may have two different register sets:

- Coprocessor general registers
- Coprocessor control registers

Each set contains up to 32 registers. Coprocessor computational instructions may use the registers in either set.

4.1.5.2 System Control Coprocessor 0 (CP0)

The system controller for all MIPS processors is implemented as coprocessor 0 (CP0²), the **System Control Coprocessor**. It provides the processor control, memory management, and exception handling functions.

4.1.5.3 Floating Point Coprocessor 1 (CP1)

If a system includes a **Floating Point Unit**, it is implemented as coprocessor 1 (CP1). Details of the FPU instructions are documented in [“Overview of the FPU Instruction Set” on page 68](#).

Coprocessor instructions are divided into two main groups:

- Load and store instructions (move to and from coprocessor), which are reserved in the main *opcode* space
- Coprocessor-specific operations, which are defined entirely by the coprocessor

4.1.5.4 Coprocessor Load and Store Instructions

Explicit load and store instructions are not defined for CP0; for CP0 only, the move to and from coprocessor instructions must be used to write and read the CP0 registers. The loads and stores for the remaining coprocessors are summarized in [“Coprocessor Loads and Stores” on page 56](#).

4.1.6 CPU Instruction Restrictions

Most 32-bit integer CPU instructions (aside from shifts) require properly sign-extended 32-bit integer operands for well-defined behavior.

2. CP0 instructions use the COP0 opcode, and as such are differentiated from the CP0 designation in this book.

Overview of the FPU Instruction Set

This chapter describes the instruction set architecture (ISA) for the floating point unit (FPU) in the microMIPS64 architecture. In the MIPS architecture, the FPU is implemented via Coprocessor 1 and Coprocessor 3, an optional processor implementing IEEE Standard 754¹ floating point operations. The FPU also provides a few additional operations not defined by the IEEE standard.

This chapter provides an overview of the following FPU architectural details:

- “Binary Compatibility” on page 68
- “Enabling the Floating Point Coprocessor” on page 69
- “IEEE Standard 754” on page 69
- “FPU Data Types” on page 69
- “Floating Point Register Types” on page 74
- “Floating Point Control Registers (FCRs)” on page 77
- “Formats of Values Used in FP Registers” on page 87
- “FPU Exceptions” on page 87
- “FPU Instructions” on page 92
- “Valid Operands for FPU Instructions” on page 98
- “FPU Instruction Formats” on page 100

The FPU instruction set is summarized by functional group. Each instruction is also described individually in alphabetical order in Volume II.

5.1 Binary Compatibility

In addition to an Instruction Set Architecture, the MIPS architecture definition includes processing resources such as the set of coprocessor general registers. A 32-bit CPU may include a full 64-bit coprocessor, including a floating point unit which implements the same mode bit to select 32-bit or 64-bit FPU register model. As of Release 5 of the

1. In this chapter, references to “IEEE standard” and “IEEE Standard 754” refer to IEEE Standard 754-1985, “IEEE Standard for Binary Floating Point Arithmetic.” For more information about this standard, see the IEEE web page at <http://grouper.ieee.org/groups/754/>.

Architecture, if floating point is implemented then FR=1 is required. I.e. the 64-bit FPU, with the FR=1 64-bit FPU register model, is required. The FR=0 32-bit FPU register model continues to be required.

Any processor implementing microMIPS64 can also run microMIPS32 binary programs, built for the same, or a lower release of the Architecture, without change.

5.2 Enabling the Floating Point Coprocessor

Enabling the Floating Point Coprocessor is done by enabling Coprocessor 1, and is a privileged operation provided by the System Control Coprocessor. If Coprocessor 1 is not enabled, an attempt to execute a floating point instruction causes a Coprocessor Unusable exception. Every system environment either enables the FPU automatically or provides a means for an application to request that it is enabled.

5.3 IEEE Standard 754

IEEE Standard 754 defines the following:

- Floating point data types
- The basic arithmetic, comparison, and conversion operations
- A computational model

The IEEE standard does not define specific processing resources nor does it define an instruction set.

The MIPS architecture includes non-IEEE FPU control and arithmetic operations (multiply-add, reciprocal, and reciprocal square root) which may not supply results that match the IEEE precision rules.

5.4 FPU Data Types

The FPU provides both floating point and fixed point data types, which are described in the next two sections.

- The single and double precision floating point data types are those specified by the IEEE standard.
- The fixed point types are signed integers provided by the CPU architecture.

5.4.1 Floating Point Formats

The following three floating point formats are provided by the FPU:

- 32-bit **single precision** floating point (type *S*, shown in [Figure 5.1](#))
- 64-bit **double precision** floating point (type *D*, shown in [Figure 5.2](#))
- 64-bit **paired single** floating point, combining two single precision data types (Type *PS*, shown in [Figure 5.3](#))

The floating point data types represent numeric values as well as other special entities, such as the following:

- Two infinities, $+\infty$ and $-\infty$

- Signaling non-numbers (SNaNs)
- Quiet non-numbers (QNaNs)s
- Numbers of the form: $(-1)^s 2^E b_0.b_1 b_2..b_{p-1}$, where
 - $s=0$ or 1
 - E =any integer between E_{min} and E_{max} , inclusive
 - $b_i=0$ or 1 (the high bit, b_0 , is to the left of the binary point)
 - p is the signed-magnitude precision




Table 5.1 Parameters of Floating Point Data Types

Parameter	Single (or each half of Paired Single)	Double
Bits of mantissa precision, p	24	53
Maximum exponent, E_{max}	+127	+1023
Minimum exponent, E_{min}	-126	-1022
Exponent <i>bias</i>	+127	+1023
Bits in exponent field, e	8	11
Representation of b_0 integer bit	hidden	hidden
Bits in fraction field, f	23	52
Total format width in bits	32	64

The single and double floating point data types are composed of three fields—*sign*, *exponent*, *fraction*—whose sizes are listed in [Table 5.1](#).

Layouts of these fields are shown in [Figures 5.1](#), [5.2](#), and [5.3](#) below. The fields are

- 1-bit sign, s
- Biased exponent, $e=E + bias$
- Binary fraction, $f=.b_1 b_2..b_{p-1}$ (the b_0 bit is not recorded)

Figure 5.1 Single-Precisions Floating Point Format (S)**Figure 5.2 Double-Precisions Floating Point Format (D)****Figure 5.3 Paired Single Floating Point Format (PS)**

Values are encoded in the specified format by using unbiased exponent, fraction, and sign values listed in [Table 5.2](#). The high-order bit of the *Fraction* field, identified as b_1 , is also important for NaNs.

Table 5.2 Value of Single or Double Floating Point DataType Encoding

Unbiased E	f	s	b_1	Value V	Type of Value	Typical Single Bit Pattern ¹	Typical Double Bit Pattern ¹
$E_{max} + 1$	$\neq 0$		1	SNaN	Signaling NaN ($FIR_{Has2008}=0$ or $FCSR_{NAN2008}=0$)	0x7fffffff	0x7fffffff ffffffff
			0	QNaN	Quiet NaN ($FIR_{Has2008}=0$ or $FCSR_{NAN2008}=0$)	0x7fbfffff	0x7ff7ffff ffffffff
$E_{max} + 1$	$\neq 0$		0	SNaN	Signaling NaN ($FCSR_{NAN2008}=1$)	0x7fbfffff	0x7ff7ffff ffffffff
			1	QNaN	Quiet NaN ($FCSR_{NAN2008}=1$)	0x7fffffff	0x7fffffff ffffffff
$E_{max} + 1$	0	1		$-\infty$	minus infinity	0xff800000	0xfff00000 00000000
		0		$+\infty$	plus infinity	0x7f800000	0x7ff00000 00000000
E_{max} to E_{min}		1		$-(2^E)(1.f)$	negative normalized number	0x80800000 through 0xff7fffff	0x80100000 00000000 through 0xffefffff ffffffff
		0		$+(2^E)(1.f)$	positive normalized number	0x00800000 through 0x7f7fffff	0x00100000 00000000 through 0x7fefffff ffffffff
$E_{min} - 1$	$\neq 0$	1		$-(2^{E_{min}})(0.f)$	negative denormalized number	0x807fffff	0x800fffff ffffffff
		0		$+(2^{E_{min}})(0.f)$	positive denormalized number	0x007fffff	0x000fffff ffffffff

Table 5.2 Value of Single or Double Floating Point Data Type Encoding (Continued)

Unbiased E	f	s	b ₁	Value V	Type of Value	Typical Single Bit Pattern ¹	Typical Double Bit Pattern ¹
$E_{min} - 1$	0	1		- 0	negative zero	0x80000000	0x80000000 00000000
		0		+ 0	positive zero	0x00000000	0x00000000 00000000

1. The "Typical" nature of the bit patterns for the NaN and denormalized values reflects the fact that the sign may have either value (NaN) and the fact that the fraction field may have any non-zero value (both). As such, the bit patterns shown are one value in a class of potential values that represent these special values.

5.4.1.1 Normalized and Denormalized Numbers

For single and double data types, each representable nonzero numerical value has just one encoding; numbers are kept in normalized form. The high-order bit of the p -bit mantissa, which lies to the left of the binary point, is “hidden,” and not recorded in the *Fraction* field. The encoding rules permit the value of this bit to be determined by looking at the value of the exponent. When the unbiased exponent is in the range E_{min} to E_{max} , inclusive, the number is normalized and the hidden bit must be 1. If the numeric value cannot be normalized because the exponent would be less than E_{min} , then the representation is denormalized and the encoded number has an exponent of $E_{min}-1$ and the hidden bit has the value 0. Plus and minus zero are special cases that are not regarded as denormalized values.

5.4.1.2 Reserved Operand Values—Infinity and NaN

A floating point operation can signal IEEE exception conditions, such as those caused by uninitialized variables, violations of mathematical rules, or results that cannot be represented. If a program does not choose to trap IEEE exception conditions, a computation that encounters these conditions proceeds without trapping but generates a result indicating that an exceptional condition arose during the computation. To permit this, each floating point format defines representations, listed in Table 5.2, for plus infinity ($+\infty$), minus infinity ($-\infty$), quiet non-numbers (QNaN), and signaling non-numbers (SNaN).

5.4.1.3 Infinity and Beyond

Infinity represents a number with magnitude too large to be represented in the format; in essence it exists to represent a magnitude overflow during a computation. A correctly signed ∞ is generated as the default result in division by zero and some cases of overflow; details are given in the IEEE exception condition described in 5.8.1 “Exception Conditions” on page 88.

Once created as a default result, ∞ can become an operand in a subsequent operation. The infinities are interpreted such that $-\infty < (\text{every finite number}) < +\infty$. Arithmetic with ∞ is the limiting case of real arithmetic with operands of arbitrarily large magnitude, when such limits exist. In these cases, arithmetic on ∞ is regarded as exact and exception conditions do not arise. The out-of-range indication represented by ∞ is propagated through subsequent computations. For some cases there is no meaningful limiting case in real arithmetic for operands of ∞ , and these cases raise the Invalid Operation exception condition (see “Invalid Operation Exception” on page 89).

5.4.1.4 Signalling Non-Number (SNaN)

SNaN operands cause the Invalid Operation exception for arithmetic operations. SNaNs are useful values to put in uninitialized variables. An SNaN is never produced as a result value.

IEEE Standard 754 states that “Whether copying a signaling NaN without a change of format signals the Invalid Operation exception is the implementor’s option.” The MIPS architecture has chosen to make the formatted operand move instructions (MOV.fmt MOVt.fmt MOVF.fmt MOVN.fmt MOVZ.fmt) non-arithmetic and they do not signal IEEE 754 exceptions.

5.4.1.5 Quiet Non-Number (QNaN)

QNaNs are intended to afford retrospective diagnostic information inherited from invalid or unavailable data and results. Propagation of the diagnostic information requires information contained in a QNaN to be preserved through arithmetic operations and floating point format conversions.

QNaN operands do not cause arithmetic operations to signal an exception. When a floating point result is to be delivered, a QNaN operand causes an arithmetic operation to supply a QNaN result. When possible, this QNaN result is one of the operand QNaN values. QNaNs do have effects similar to SNaNs on operations that do not deliver a floating point result—specifically, comparisons. (For more information, see the detailed description of the floating point compare instruction, C.cond.fmt.)

When certain invalid operations not involving QNaN operands are performed but do not trap (because the trap is not enabled), a new QNaN value is created. Table 5.3 shows the QNaN value generated when no input operand QNaN value can be copied. The values listed for the fixed point formats are the values supplied to satisfy the IEEE standard when a QNaN or infinite floating point value is converted to fixed point. There is no other feature of the architecture that detects or makes use of these “integer QNaN” values. The $FCSR_{NAN2008}=1$ “integer QNaN” values were chosen to match the requirements of the Java and Fortran programming languages.

Table 5.3 Value Supplied When a New Quiet NaN Is Created

Format	New QNaN value ($FIR_{Has2008} = 0$ or $FCSR_{NAN2008} = 0$)	New QNaN value ($FCSR_{NAN2008} = 1$)
Single floating point	0x7fbf ffff	0x7fc0 0000
Double floating point	0x7ff7 ffff ffff ffff	0x7ff8 0000 0000 0000
Word fixed point (result from converting any FP number too big to represent as a 32-bit positive integer)	0x7fff ffff	0x7fff ffff
Word fixed point (result from converting any FP NaN)	0x7fff ffff	0x0000 0000
Word fixed point (result from converting any FP number too small to represent as a 32-bit negative integer)	0x7fff ffff	0x8000 0000
Longword fixed point (result from converting any FP number too big to represent as a 64-bit positive integer)	0x7fff ffff ffff ffff	0x7fff ffff ffff ffff
Longword fixed point (result from converting any FP NaN)	0x7fff ffff ffff ffff	0x0000 0000 0000 0000
Longword fixed point(result from con- verting any FP number too small to represent as a 64-bit negative integer)	0x7fff ffff ffff ffff	0x8000 0000 0000 0000

If a CPU implements passing an input NAN operand to the output of an instruction in hardware (instead of taking an Unimplemented FP exception) and $FCSR_{NAN2008}=1$, the mantissa portion of the input NAN operand is preserved as much as possible:

- If the chosen input is a QNAN, the entire mantissa is passed to the output without change.
- If the chosen input is a SNAN, the only change is to set the leftmost/most-significant mantissa bit.

5.4.1.6 Paired Single Exceptions

Exception conditions that arise while executing the two halves of a floating point vector operation are ORed together, and the instruction is treated as having caused all the exceptional conditions arising from both operations. The hardware makes no effort to determine which of the two operations encountered the exceptional condition.

5.4.1.7 Paired Single Condition Codes

The c.cond.PS instruction compares the upper and lower halves of FPR *fs* and FPR *ft* independently and writes the results into condition codes CC +1 and CC respectively. The CC number must be even. If the number is not even the operation of the instruction is **UNPREDICTABLE**.

5.4.2 Fixed Point Formats

The FPU provides two fixed point data types:

- 32-bit **Word** fixed point (type *W*), shown in [Figure 5.4](#)
- 64-bit **Longword** fixed point (type *L*), shown in [Figure 5.5](#)

The fixed point values are held in the 2's complement format used for signed integers in the CPU. Unsigned fixed point data types are not provided by the architecture; application software may synthesize computations for unsigned integers from the existing instructions and data types.

Figure 5.4 Word Fixed Point Format (W)





Figure 5.5 Longword Fixed Point Format (L)



5.5 Floating Point Register Types

This section describes the organization and use of the two types of FPU register sets:

Floating Point registers (FPRs) are 32 or 64 bits wide. A 32-bit floating point unit contains 32 32-bit FPRs, each of which is capable of storing a 32-bit data type. Double-precision (type D) data types are stored in even-odd pairs of FPRs, and the long-integer (type L) and paired single (type PS) data types are not supported. A 64-bit floating point unit contains 32 64-bit FPRs, each of which is capable of storing any data type. For compatibility with 32-bit FPUs, the FR bit in the CP0 *Status* register is used by processor that supports a 64-bit FPU to configure the FPU in a mode in which the FPRs are treated as 32 32-bit registers, each of which is capable of storing only 32-bit data types. In this mode, the double-precision floating point (type D) data type is stored in even-odd pairs of FPRs, and the long-integer (type L) and paired single (type PS) data types are not supported.

- These registers transfer binary data between the FPU and the system, and are also used to hold formatted FPU operand values. Refer to *Volume III, The MIPS Privileged Architecture Manual*, for more information on the CP0 Registers.
- *Floating Point Control* registers (FCRs), which are 32 bits wide. There are five FPU control registers, used to identify and control the FPU. These registers are indicated by the *fs* field of the instruction word. Three of these registers, *FCCR*, *FEXR*, and *FENR*, select subsets of the floating point *Control/Status* register, the *FCSR*.

5.5.1 FPU Register Models

There are separate FPU register models:

- 32 32-bit registers, with D-format values stored in even-odd pairs of registers.
- 32 64-bit registers, with all formats supported in a register.

The register models are made available as a mode selection through the **FR Bit of the CP0 Status Register**.

If the value of FR bit is changed, the contents of the FPRs becomes **UNPREDICTABLE**. For some implementations, it might be necessary for software to re-initialize the FPRs.

5.5.2 Binary Data Transfers (32-Bit and 64-Bit)

The data transfer instructions move words and doublewords between the FPU FPRs and the remainder of the system. The operations of the word and doubleword load and move-to instructions are shown in [Figure 5.6](#) and [Figure 5.7](#).

The store and move-from instructions operate in reverse, reading data from the location which the corresponding load or move-to instruction wrote.

Figure 5.6 FPU Word Load and Move-to Operations








Figure 5.7 FPU Doubleword Load and Move-to Operations



5.5.3 FPRs and Formatted Operand Layout

FPU instructions that operate on formatted operand values specify the **floating point register** (FPR) that holds the value. Operands that are only 32 bits wide (*W* and *S* formats), use only half the space in a 64-bit FPR.

The FPR organization and the way that operand data is stored in them is shown in Figures 5.8, 5.9 and 5.10.

Figure 5.8 Single Floating Point or Word Fixed Point Operand in an FPR**Figure 5.9 Double Floating Point or Longword Fixed Point Operand in an FPR****Figure 5.10 Paired-Single Floating Point Operand in an FPR**

5.6 Floating Point Control Registers (FCRs)

Access to the Floating Point Control Registers is not privileged; they can be accessed by any program that can execute floating point instructions. The FCRs can be accessed via the CTC1 and CFC1 instructions.

The microMIPS64 Architecture supports the following floating point *Control* registers (*FCRs*):

- *FIR*, *FP Implementation and Revision* register
- *FCSR*, *FP Control/Status* register (used to be known as *FCR31*).

The following FCRs are “aliases”, accessing subsets of the *FCSR* through CTC1 and CFC1 instructions.

- *FCCR*, *FP Condition Codes* register
- *FEXR*, *FP Exceptions* register
- *FENR*, *FP Enables* register

Release 5 of the Architecture adds additional “alias” FCRs:

- *UFR* and *UNFR* for user-mode access to *Status_{FR}*. See CTC1, CFC1 for Release 5 related changes.

5.6.1 Floating Point Implementation Register (FIR, CP1 Control Register 0)

Compliance Level: *Required* if floating point is implemented

The Floating Point Implementation Register (*FIR*) is a 32-bit read-only register that contains information identifying the capabilities of the floating point unit, the floating point processor identification, and the revision level of the floating point unit. [Figure 5.11](#) shows the format of the *FIR* register; [Table 5.4](#) describes the *FIR* register fields.

Figure 5.11 FIR Register Format

31	29	28	27	24	23	22	21	20	19	18	17	16	15	8	7	0
0 000	UFRP	Impl	Has 2008	F64	L	W	3D	PS	D	S	ProcessorID					Revision

Table 5.4 FIR Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance						
Name	Bits										
0	31:28	Reserved for future use; reads as zero	0	0	Reserved						
UFRP	28	See Release 5 definition of CFC1 and CTC1. <div><table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>User mode FR switching instructions not supported.</td></tr><tr><td>1</td><td>User mode FR switching instructions supported.</td></tr></table></div>	Encoding	Meaning	0	User mode FR switching instructions not supported.	1	User mode FR switching instructions supported.	R	Preset	Optional (Release 5)
Encoding	Meaning										
0	User mode FR switching instructions not supported.										
1	User mode FR switching instructions supported.										
Impl	27..24	These bits are implementation dependent and are not defined by the architecture, other than the fact that they are read-only. This bits are explicitly not intended to be used for mode control functions.	R	Preset	Optional						
Has2008	23	Indicates that one or more IEEE-754-2008 features are implemented. If this bit is set, the ABS2008 and NAN2008 fields within the FCSR register also exist.	R	Preset by hardware	Optional as of Release 3. Required as of Release 5.						
F64	22	Indicates that the floating point unit has registers and data paths that are 64-bits wide.. <div><table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>FPU is 32 bits</td></tr><tr><td>1</td><td>FPU is 64 bits</td></tr></table></div>	Encoding	Meaning	0	FPU is 32 bits	1	FPU is 64 bits	R	Preset by hardware	Required
Encoding	Meaning										
0	FPU is 32 bits										
1	FPU is 64 bits										
L	21	Indicates that the longword fixed point (L) data type and instructions are implemented: <div><table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>L fixed point not implemented</td></tr><tr><td>1</td><td>L fixed point implemented</td></tr></table></div>	Encoding	Meaning	0	L fixed point not implemented	1	L fixed point implemented	R	Preset by hardware	Required
Encoding	Meaning										
0	L fixed point not implemented										
1	L fixed point implemented										

Table 5.4 FIR Register Field Descriptions (Continued)

Fields		Description	Read/ Write	Reset State	Compliance												
Name	Bits																
W	20	Indicates that the word fixed point (W) data type and instructions are implemented: <table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>W fixed point not implemented</td></tr><tr><td>1</td><td>W fixed point implemented</td></tr></table>	Encoding	Meaning	0	W fixed point not implemented	1	W fixed point implemented	R	Preset by hardware	Required						
Encoding	Meaning																
0	W fixed point not implemented																
1	W fixed point implemented																
3D	19	The MIPS-3D ASE is supported on any processors with a 64-bit floating point unit, and this bit indicates that the MIPS-3D ASE is implemented: <table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>MIPS-3D ASE not implemented</td></tr><tr><td>1</td><td>MIPS-3D ASE implemented</td></tr></table> Indicates that the MIPS-3D ASE is implemented: <table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>MIPS-3D ASE not implemented</td></tr><tr><td>1</td><td>MIPS-3D ASE implemented</td></tr></table>	Encoding	Meaning	0	MIPS-3D ASE not implemented	1	MIPS-3D ASE implemented	Encoding	Meaning	0	MIPS-3D ASE not implemented	1	MIPS-3D ASE implemented	R	Preset by hardware	Required
Encoding	Meaning																
0	MIPS-3D ASE not implemented																
1	MIPS-3D ASE implemented																
Encoding	Meaning																
0	MIPS-3D ASE not implemented																
1	MIPS-3D ASE implemented																
PS	18	Indicates that the paired single floating point data type is implemented: <table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>PS floating point not implemented</td></tr><tr><td>1</td><td>PS floating point implemented</td></tr></table>	Encoding	Meaning	0	PS floating point not implemented	1	PS floating point implemented	R	Preset by hardware	Required						
Encoding	Meaning																
0	PS floating point not implemented																
1	PS floating point implemented																
D	17	Indicates that the double-precision (D) floating point data type and instructions are implemented: <table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>D floating point not implemented</td></tr><tr><td>1</td><td>D floating point implemented</td></tr></table>	Encoding	Meaning	0	D floating point not implemented	1	D floating point implemented	R	Preset by hardware	Required						
Encoding	Meaning																
0	D floating point not implemented																
1	D floating point implemented																
S	16	Indicates that the single-precision (S) floating point data type and instructions are implemented: <table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>S floating point not implemented</td></tr><tr><td>1</td><td>S floating point implemented</td></tr></table>	Encoding	Meaning	0	S floating point not implemented	1	S floating point implemented	R	Preset by hardware	Required						
Encoding	Meaning																
0	S floating point not implemented																
1	S floating point implemented																
ProcessorID	15:8	Identifies the floating point processor.	R	Preset by hardware	Required												

Table 5.4 FIR Register Field Descriptions (Continued)

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
Revision	7:0	Specifies the revision number of the floating point unit. This field allows software to distinguish between one revision and another of the same floating point processor type. If this field is not implemented, it must read as zero.	R	Preset by hardware	Optional

5.6.2 User Floating point Register mode control (UFR, CP1 Control Register 1)

Compliance Level: *Optional* in MIPS64r5 if floating point is implemented and user-mode FR switching is supported

The UFR register allows user-mode to clear $Status_{FR}$ by executing a CTC1 to UFR with GPR[0] as input, and read $Status_{FR}$ by executing a CFC1 to UFR. CTC1 to UFR with any other input register is required to produce a Reserved Instruction Exception. User-mode software can determine presence of this feature from FIR_{UFRP} .

Per the definition of the CTC1 instruction, writing any value other than 0 obtained from integer GPR[0] to UFR using the CTC1 instruction is UNPREDICTABLE. To set UFR.FR / $Status_{FR}$ use CTC1 to the UNFR FCR alias.

Figure 5.12 UFR Register Format




Table 5.5 UFR Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
0	31:1	Must be written as zero; returns zero on read	0	0	Reserved
FR	0	User-mode access to $Status_{FR}$.	R/W0g ¹	Undefined See reset state of $Status_{FR}$	Optional (Release 5)

1. R/W0g definition: UFR can read as 0 or 1, but can only be written with the zero from GPR[0], which clears $Status_{FR}$. Using CTC1 to write UFR with any value or GPR other than GPR[0] is UNPREDICTABLE.

5.6.3 User Negated Floating point Register mode control (UNFR, CP1 Control Register 4)

Compliance Level: *Optional* in MIPS64r5 if floating point is implemented and user-mode FR switching is supported

The UNFR register allows user-mode to set $Status_{FR}$ by executing a CTC1 to UNFR with GPR[0] as input. CTC1 to UNFR with any other input register is required to produce a Reserved Instruction Exception. User-mode software can determine presence of this feature from FIR_{UFRP} .

Figure 5.13 UNFR Register Format




Table 5.6 UNFR Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
0	31:1	Must be written as zero.	0	n.a. ²	Reserved
NFR	0	User-mode inverted write to <i>Status_{FR}</i> .	W0 ¹	n.a. ²	Optional (Release 5)

1. W0g definition: UNFR can only be written with the zero from GPR[0], which sets *Status_{FR}*. Using CFC1 to read UNFR, or using CTC1 to write UNFR with any value or GPR other than GPR[0] is UNPREDICTABLE. UNFR's "state" can be inferred by reading *Status_{FR}*, e.g. via UFR.
2. UNFR's reset state is "n.a." (not applicable), since UNFR is not readable state.

The UNFR pseudo-control-register alias is a convenience, allowing CTC1 R0, UNFR to be used to set UFR/*Status_{FR}* without requiring a GPR to hold the value such as 1 to be written. Because reading UNFR would be redundant with reading UFR, UNFR is write-only; attempting to read UNFR via CFC1 is UNPREDICTABLE, per the definition of the CFC1 instruction. Writing any value other than 0 obtained from integer GPR R0 to UNFR using the CTC1 instruction is similarly UNPREDICTABLE from software's point of view, and is required to produce a Reserved Instruction Exception in R5.03 implementations.


5.6.4 Floating Point Control and Status Register (FCSR, CP1 Control Register 31)

Compliance Level: *Required* if floating point is implemented.

The Floating Point Control and Status Register (*FCSR*) is a 32-bit register that controls the operation of the floating point unit, and shows the following status information:

- selects the default rounding mode for FPU arithmetic operations
- selectively enables traps of FPU exception conditions
- controls some denormalized number handling options
- reports any IEEE exceptions that arose during the most recently executed instruction
- reports IEEE exceptions that arose, cumulatively, in completed instructions
- indicates the condition code result of FP compare instructions

Access to *FCSR* is not privileged; it can be read or written by any program that has access to the floating point unit (via the coprocessor enables in the *Status* register). Figure 5.14 shows the format of the *FCSR* register; Table 5.7 describes the *FCSR* register fields.

Figure 5.14 FCSR Register Format**Table 5.7 FCSR Register Field Descriptions**

Fields		Description	Read/ Write	Reset State	Compliance						
Name	Bits										
FCC	31:25, 23	Floating point condition codes. These bits record the result of floating point compares and are tested for floating point conditional branches and conditional moves. The FCC bit to use is specified in the compare, branch, or conditional move instruction. For backward compatibility with previous MIPS ISAs, the FCC bits are separated into two, non-contiguous fields.	R/W	Undefined	Required						
FS	24	<div>Flush to Zero (Flush Subnormals). See sections 5.8.1.3 “Underflow Exception” on page 90 and 5.8.1.4 “Alternate Flush to Zero Underflow Handling” on page 91.</div> <table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>Input subnormal values and tiny non-zero results are not altered. Unimplemented Operation Exception may be signaled as needed.</td></tr><tr><td>1</td><td>When FS is one, subnormal results are flushed to zero. The Unimplemented Operation Exception is NOT signalled for this reason. <div>Every tiny non-zero result is replaced with zero of the same sign. Prior to Release 5 it is implementation dependent whether subnormal operand values are flushed to zero before the operation is carried out. As of Release 5 every input subnormal value is replaced with zero of the same sign.</div></td></tr></table>	Encoding	Meaning	0	Input subnormal values and tiny non-zero results are not altered. Unimplemented Operation Exception may be signaled as needed.	1	When FS is one, subnormal results are flushed to zero. The Unimplemented Operation Exception is NOT signalled for this reason. <div>Every tiny non-zero result is replaced with zero of the same sign. Prior to Release 5 it is implementation dependent whether subnormal operand values are flushed to zero before the operation is carried out. As of Release 5 every input subnormal value is replaced with zero of the same sign.</div>	R/W	Undefined	Required
Encoding	Meaning										
0	Input subnormal values and tiny non-zero results are not altered. Unimplemented Operation Exception may be signaled as needed.										
1	When FS is one, subnormal results are flushed to zero. The Unimplemented Operation Exception is NOT signalled for this reason. <div>Every tiny non-zero result is replaced with zero of the same sign. Prior to Release 5 it is implementation dependent whether subnormal operand values are flushed to zero before the operation is carried out. As of Release 5 every input subnormal value is replaced with zero of the same sign.</div>										
Impl	22:21	Available to control implementation dependent features of the floating point unit. If these bits are not implemented, they must be ignored on write and read as zero.	R/W	Undefined	Optional						
0	20	Reserved for future use; reads as zero.	R	Preset by hardware	Reserved						

Table 5.7 FCSR Register Field Descriptions (Continued)

Fields		Description	Read/ Write	Reset State	Compliance						
Name	Bits										
ABS2008	19	<p>ABS.fmt & NEG.fmt instructions compliant with IEEE Standard 754-2008.</p> <p>The IEEE 754-2008 standard requires that the ABS and NEG functions are non-arithmetic and accept NAN inputs without trapping.</p> <p>This fields exists if $FIR_{Has2008}$ is set.</p> <table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>ABS & NEG instructions are arithmetic and trap for NAN input. MIPS legacy behavior.</td></tr><tr><td>1</td><td>ABS & NEG instructions are non-arithmetic and accept NAN input without trapping. IEEE 754-2008 behavior</td></tr></table>	Encoding	Meaning	0	ABS & NEG instructions are arithmetic and trap for NAN input. MIPS legacy behavior.	1	ABS & NEG instructions are non-arithmetic and accept NAN input without trapping. IEEE 754-2008 behavior	R	Preset by hardware	Required as of Release 5
Encoding	Meaning										
0	ABS & NEG instructions are arithmetic and trap for NAN input. MIPS legacy behavior.										
1	ABS & NEG instructions are non-arithmetic and accept NAN input without trapping. IEEE 754-2008 behavior										
NAN2008	18	<p>Quiet and signaling NaN encodings recommended by the IEEE Standard 754-2008, i.e. a quiet NaN is encoded with the first bit of the fraction being 1 and a signaling NaN is encoded with the first bit of the fraction field being 0. MIPS legacy FPU encodes NaN values with the opposite polarity, i.e. a quiet NaN is encoded with the first bit of the fraction being 0 and a signaling NaN is encoded with the first bit of the fraction field being 1. Refer to Table 5.3 for the quiet NaN encoding values.</p> <p>This fields exists if $FIR_{Has2008}$ is set.</p> <table><tr><th>Encoding</th><th>Meaning</th></tr><tr><td>0</td><td>MIPS legacy NaN encoding</td></tr><tr><td>1</td><td>IEEE 754-2008 NaN encoding</td></tr></table>	Encoding	Meaning	0	MIPS legacy NaN encoding	1	IEEE 754-2008 NaN encoding	R	Preset by hardware	Required as of Release 5
Encoding	Meaning										
0	MIPS legacy NaN encoding										
1	IEEE 754-2008 NaN encoding										
Cause	17:12	<p>Cause bits. These bits indicate the exception conditions that arise during execution of an FPU arithmetic instruction. A bit is set to 1 if the corresponding exception condition arises during the execution of an instruction and is set to 0 otherwise. By reading the registers, the exception condition caused by the preceding FPU arithmetic instruction can be determined.</p> <p>Refer to Table 5.8 for the meaning of each bit.</p>	R/W	Undefined	Required						

Table 5.7 FCSR Register Field Descriptions (Continued)

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
Enables	11:7	Enable bits. These bits control whether or not an exception is taken when an IEEE exception condition occurs for any of the five conditions. The exception occurs when both an Enable bit and the corresponding Cause bit are set either during an FPU arithmetic operation or by moving a value to FCSR or one of its alternative representations. Note that Cause bit E has no corresponding Enable bit; the non-IEEE Unimplemented Operation exception is defined by MIPS as always enabled. Refer to Table 5.8 for the meaning of each bit.	R/W	Undefined	Required
Flags	6:2	Flag bits. This field shows any exception conditions that have occurred for completed instructions since the flag was last reset by software. When a FPU arithmetic operation raises an IEEE exception condition that does not result in a Floating Point Exception (i.e., the Enable bit was off), the corresponding bit(s) in the Flag field are set, while the others remain unchanged. Arithmetic operations that result in a Floating Point Exception (i.e., the Enable bit was on) do not update the Flag bits. This field is never reset by hardware and must be explicitly reset by software. Refer to Table 5.8 for the meaning of each bit.	R/W	Undefined	Required
RM	1:0	Rounding mode. This field indicates the rounding mode used for most floating point operations (some operations use a specific rounding mode). Refer to Table 5.9 for the meaning of the encodings of this field.	R/W	Undefined	Required.

The FCC, FS, Cause, Enables, Flags and RM fields in the *FCSR*, *FCCR*, *FEXR*, and *FENR* registers always display the correct state. That is, if a field is written via *FCCR*, the new value may be read via one of the alternate registers. Similarly, if a value is written via one of the alternate registers, the new value may be read via *FCSR*.

Table 5.8 Cause, Enable, and Flag Bit Definitions

Bit Name	Bit Meaning
E	Unimplemented Operation (this bit exists only in the Cause field)
V	Invalid Operation
Z	Divide by Zero
O	Overflow
U	Underflow
I	Inexact


Table 5.9 Rounding Mode Definitions

RM Field Encoding	Meaning
0	RN - Round to Nearest Rounds the result to the nearest representable value. When two representable values are equally near, the result is rounded to the value whose least significant bit is zero (that is, even)
1	RZ - Round Toward Zero Rounds the result to the value closest to but not greater than in magnitude than the result.
2	RP - Round Towards Plus Infinity Rounds the result to the value closest to but not less than the result.
3	RM - Round Towards Minus Infinity Rounds the result to the value closest to but not greater than the result.

5.6.5 Floating Point Condition Codes Register (FCCR, CP1 Control Register 25)

Compliance Level: *Required* if floating point is implemented.

The Floating Point Condition Codes Register (*FCCR*) is an alternative way to read and write the floating point condition code values that also appear in *FCSR*. Unlike *FCSR*, all eight FCC bits are contiguous in *FCCR*. [Figure 5.15](#) shows the format of the *FCCR* register; [Table 5.10](#) describes the *FCCR* register fields.

Figure 5.15 FCCR Register Format**Table 5.10 FCCR Register Field Descriptions**

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
0	31:8	Must be written as zero; returns zero on read	0	0	Reserved
FCC	7:0	Floating point condition code. Refer to the description of this field in the <i>FCSR</i> register.	R/W	Undefined	Required

5.6.6 Floating Point Exceptions Register (FEXR, CP1 Control Register 26)

Compliance Level: *Required* if floating point is implemented.

The Floating Point Exceptions Register (*FEXR*) is an alternative way to read and write the Cause and Flags fields that also appear in *FCSR*. [Figure 5.16](#) shows the format of the *FEXR* register; [Table 5.11](#) describes the *FEXR* register fields.

Figure 5.16 FEXR Register Format




Table 5.11 FEXR Register Field Descriptions

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
0	31:18, 11:7, 1:0	Must be written as zero; returns zero on read	0	0	Reserved
Cause	17:12	Cause bits. Refer to the description of this field in the <i>FCSR</i> register.	R/W	Undefined	Required
Flags	6:2	Flags bits. Refer to the description of this field in the <i>FCSR</i> register.	R/W	Undefined	Optional

5.6.7 Floating Point Enables Register (FENR, CP1 Control Register 28)

Compliance Level: *Required* if floating point is implemented.

The Floating Point Enables Register (*FENR*) is an alternative way to read and write the Enables, FS, and RM fields that also appear in *FCSR*. Figure 5.17 shows the format of the *FENR* register; Table 5.12 describes the *FENR* register fields.

31	12	11	10	9	8	7	6	3	2	1	0
0 0000 0000 0000 0000 0000						Enables		0 000 0		FS	RM
						V	Z	O	U	I	

Fields		Description	Read/ Write	Reset State	Compliance
Name	Bits				
0	31:12, 6:3	Must be written as zero; returns zero on read	0	0	Reserved
Enables	11:7	Enable bits. Refer to the description of this field in the <i>FCSR</i> register.	R/W	Undefined	Required
FS	2	Flush to Zero bit. Refer to the description of this field in the <i>FCSR</i> register.	R/W	Undefined	Required
RM	1:0	Rounding mode. Refer to the description of this field in the <i>FCSR</i> register.	R/W	Undefined	Required

Unlike the CPU, the FPU does not interpret the binary encoding of source operands nor produce a binary encoding of results for every operation. The value held in a floatig point register (FPR) is either *uninterpreted*, or one one of the valid numeric formats: *single*, *double*, *paired-single* floating-point, *word* and *long* fixed point.

- When a data transfer instruction writes binary data into a FPR (LWC1, LWXC1, LDC1, LDXC1, LUXC1, MTC1, MTHC1, DMTC1), then the binary value of the register is *uninterpreted*.
- A FP computational or FP register move (MOV*.fmt) instruction which produces a result of type *fnt* puts a value of type *fnt* into the result register.
- The format of the value of a FPR is unchanged when it is read by data transfer instruction (SWC1, SWXC1, SDC1, SDXC1, SUXC1, MFC1, MFHC1, DMFC1).

When an FPR with an uninterpreted value is used as a source operand by an instruction that requires a value of format *fmt*, the binary contents are interpreted as a value of format *fmt*. A FP arithmetic instruction produces a value of the expected numeric format into the destination register.

If an FPR contains a value of numeric format *fnt* and an instruction uses the FPR as source operand of different numeric format, the result of the instruction is **UNPREDICTABLE**.

This section provides the following information FPU exceptions:

- Precise exception mode
- Descriptions of the exceptions

- Non-Arithmetic Instructions

FPU exceptions are implemented in the MIPS FPU architecture with the *Cause*, *Enable*, and *Flag* fields of the *Control/Status* register. The *Flag* bits implement IEEE exception status flags, and the *Cause* and *Enable* bits control exception trapping. Each field has a bit for each of the five IEEE exception conditions and the *Cause* field has an additional exception bit, Unimplemented Operation, used to trap for software emulation assistance.

5.8.0.1 Precise Exception Mode

In precise exception mode, a trap occurs before the instruction that causes the trap, or any following instruction, can complete and write its results. If desired, the software trap handler can resume execution of the interrupted instruction stream after handling the exception.

The *Cause* field reports per-bit instruction exception conditions. The *Cause* bits are written during each floating point arithmetic operation to show any exception conditions that arise during the operation. The bit is set to 1 if the corresponding exception condition arises; otherwise it is set to 0.

A floating point trap is generated any time both a *Cause* bit and its corresponding *Enable* bit are set. This occurs either during the execution of a floating point operation or by moving a value into the *FCSR*. There is no *Enable* for Unimplemented Operation; this exception always generates a trap.

In a trap handler, exception conditions that arise during any trapped floating point operations are reported in the *Cause* field. Before returning from a floating point interrupt or exception, or before setting *Cause* bits with a move to the *FCSR*, software must first clear the enabled *Cause* bits by executing a move to *FCSR* to prevent the trap from being erroneously retaken.

User-mode programs cannot observe enabled *Cause* bits being set. If this information is required in a User-mode handler, it must be available someplace other than through the *Status* register.

If a floating point operation sets only non-enabled *Cause* bits, no trap occurs and the default result defined by the IEEE standard is stored (see Table 5.13). When a floating point operation does not trap, the program can monitor the exception conditions by reading the *Cause* field.

The *Flag* field is a cumulative report of IEEE exception conditions that arise as instructions complete; instructions that trap do not update the *Flag* bits. The *Flag* bits are set to 1 if the corresponding IEEE exception is raised, otherwise the bits are unchanged. There is no *Flag* bit for the MIPS Unimplemented Operation exception. The *Flag* bits are never cleared as a side effect of floating point operations, but may be set or cleared by moving a new value into the *FCSR*.

Addressing exceptions are precise.

5.8.1 Exception Conditions

The following five exception conditions defined by the IEEE standard are described in this section:

- Invalid Operation Exception
- Division By Zero Exception
- Underflow Exception
- Overflow Exception

- [Inexact Exception](#)

This section also describes a MIPS-specific exception condition, **Unimplemented Operation**, that is used to signal a need for software emulation of an instruction. Normally an IEEE arithmetic operation can cause only one exception condition; the only case in which two exceptions can occur at the same time are Inexact With Overflow and Inexact With Underflow.

At the program's direction, an IEEE exception condition can either cause a trap or not cause a trap. The IEEE standard specifies the result to be delivered in case the exception is not enabled and no trap is taken. The MIPS architecture supplies these results whenever the exception condition does not result in a precise trap (that is, no trap or an imprecise trap). The default action taken depends on the type of exception condition, and in the case of the Overflow, the current rounding mode. The default results are summarized in [Table 5.13](#).

Table 5.13 Default Result for IEEE Exceptions Not Trapped Precisely

Bit	Description	Default Action
V	Invalid Operation	Supplies a quiet NaN.
Z	Divide by zero	Supplies a properly signed infinity.
U	Underflow	Supplies a rounded result.
I	Inexact	Supplies a rounded result. If caused by an overflow without the overflow trap enabled, supplies the overflowed result.
O	Overflow	Depends on the rounding mode, as shown below.
	0 (RN)	Supplies an infinity with the sign of the intermediate result.
	1 (RZ)	Supplies the format's largest finite number with the sign of the intermediate result.
	2 (RP)	For positive overflow values, supplies positive infinity. For negative overflow values, supplies the format's most negative finite number.
	3 (RM)	For positive overflow values, supplies the format's largest finite number. For negative overflow values, supplies minus infinity.

5.8.1.1 Invalid Operation Exception

The Invalid Operation exception is signaled if one or both of the operands are invalid for the operation to be performed. The result, when the exception condition occurs without a precise trap, is a quiet NaN.

These are invalid operations:

- One or both operands are a signaling NaN (except for the non-arithmetic MOV.fmt, MOVT.fmt, MOVF.fmt, MOVN.fmt, and MOVZ.fmt instructions).
- Addition or subtraction: magnitude subtraction of infinities, such as $(+\infty) + (-\infty)$ or $(-\infty) - (-\infty)$.
- Multiplication: $0 \times \infty$, with any signs.
- Division: $0/0$ or ∞/∞ , with any signs.
- Square root: An operand of less than 0 (-0 is a valid operand value).
- Conversion of a floating point number to a fixed point format when either an overflow or an operand value of infinity or NaN precludes a faithful representation in that format.

- Some comparison operations in which one or both of the operands is a QNaN value. (The detailed definition of the compare instruction, C.cond.fmt, in Volume II has tables showing the comparisons that do and do not signal the exception.)

5.8.1.2 Division By Zero Exception

An implemented divide operation signals a Division By Zero exception if the divisor is zero and the dividend is a finite nonzero number. The result, when no precise trap occurs, is a correctly signed infinity. Divisions (0/0) and (∞ /0) do not cause the Division By Zero exception. The result of (0/0) is an Invalid Operation exception. The result of (∞ /0) is a correctly signed infinity.

5.8.1.3 Underflow Exception

This section describes IEEE standard compliant underflow exception handling, desired when FCSR.FS=0. Some implementations may require software assistance to accomplish this, via the Unimplemented Operation Exception handler. See the next section, 5.8.1.4, for [Alternate Flush to Zero Underflow Handling](#), obtained by setting FCSR.FS=1, which may be faster on some implementations.

Two related events contribute to underflow:

- **Tininess:** the creation of a tiny nonzero result between $\pm 2^{E_{min}}$ which, because it is tiny, may cause some other exception later such as overflow on division
- **Loss of accuracy:** the extraordinary loss of accuracy during the approximation of such tiny numbers by denormalized numbers

Tininess: The IEEE standard allows choices in detecting these events, but requires that they be detected in the same manner for all operations. The IEEE standard specifies that “tininess” may be detected at either of these times:

- *After rounding*, when a nonzero result computed as though the exponent range were unbounded would lie strictly between $\pm 2^{E_{min}}$
- *Before rounding*, when a nonzero result computed as though both the exponent range and the precision were unbounded would lie strictly between $\pm 2^{E_{min}}$

The MIPS architecture specifies that tininess be detected after rounding.

Loss of Accuracy: The IEEE standard specifies that loss of accuracy may be detected as a result of either of these conditions:

- *Denormalization loss*, when the delivered result differs from what would have been computed if the exponent range were unbounded
- *Inexact result*, when the delivered result differs from what would have been computed if both the exponent range and precision were unbounded

The MIPS architecture specifies that loss of accuracy is detected as inexact result.

Signalling an Underflow: When an underflow trap is not enabled, underflow is signaled only when both tininess and loss of accuracy have been detected. The delivered result might be zero, denormalized, or $\pm 2^{E_{min}}$.

When an underflow trap is enabled (through the *FCSR Enable* field bit), underflow is signaled when tininess is detected regardless of loss of accuracy.

5.8.1.4 Alternate Flush to Zero Underflow Handling

Previous section 5.8.1.3 “Underflow Exception” describes IEEE standard compliant underflow exception handling, desired when *FCSR.FS*=0. The current section describes [Alternate Flush to Zero Underflow Handling](#), obtained by setting *FCSR.FS*=1, which never requires the Unimplemented Operation Exception handler to handle subnormal results, and which may be faster on some implementations even if software exception handler assistance is not required.

When the *FCSR.FS* is set:

Results: Every tiny non-zero result is replaced with zero of the same sign.

Inputs: Prior to Release 5 it is implementation dependent whether subnormal operand values are flushed to zero before the operation is carried out. As of Release 5 every input subnormal value is replaced with zero of the same sign.

Exceptions: Because the *FCSR.FS* bit flushes subnormal results to zero, the Unimplemented Operation Exception will never be produced for this reason. All the other floating point exceptions are signaled according to the new values of the operands or the results. In addition, when the *FCSR.FS* bit is set:

- Tiny non-zero results are detected before rounding². Flushing of tiny non-zero results causes Inexact and Underflow Exceptions to be signaled.
- Flushing of subnormal input operands in all instructions except comparisons causes Inexact Exception to be signaled.
- For floating-point comparisons, the Inexact Exception is not signaled when subnormal input operands are flushed.
- Inputs to non-arithmetic floating-point instructions are never flushed.

Should the alternate exception handling attributes of the IEEE Standard for Floating-Point Arithmetic 754TM-2008, Section 8 be desired, the *FCSR.FS* bit should be zero, the Underflow Exception should be enabled and a trap handler should be provided to carry out the execution of the alternate exception handling attributes.

5.8.1.5 Overflow Exception

An Overflow exception is signaled when the magnitude of a rounded floating point result, were the exponent range unbounded, is larger than the destination format's largest finite number.

When no precise trap occurs, the result is determined by the rounding mode and the sign of the intermediate result.

5.8.1.6 Inexact Exception

An Inexact exception is signaled if one of the following occurs:

- The rounded result of an operation is not exact

2. Tiny non-zero results that would have been normal after rounding are flushed to zero.

- The rounded result of an operation overflows without an overflow trap

5.8.1.7 Unimplemented Operation Exception

The Unimplemented Operation exception is a MIPS defined exception that provides software emulation support. This exception is not IEEE-compliant.

The MIPS architecture is designed so that a combination of hardware and software may be used to implement the architecture. Operations that are not fully supported in hardware cause an Unimplemented Operation exception so that software may perform the operation.

There is no *Enable* bit for this condition; it always causes a trap. After the appropriate emulation or other operation is done in a software exception handler, the original instruction stream can be continued.

5.8.1.8 Non-Arithmetic Instructions

Some FPU conversion and FPU Formatted Operand-Value Move instructions (see next section) do not perform floating-point arithmetic operations on their input operands. For that reason, such instructions do not generate IEEE arithmetic exceptions. These instructions include MOV.fmt, MOVF.fmt, MOVT.fmt, MOVZ.fmt, MOVN.fmt, PLL.PS, PLU.PS, PUL.PS, PUU.PS, CVT.S.PU, CVT.PS.S, CVT.S.PL.

5.9 FPU Instructions

The FPU instructions comprise the following functional groups:

- [Data Transfer Instructions](#)
- [Arithmetic Instructions](#)
- [Conversion Instructions](#)
- [Formatted Operand-Value Move Instructions](#)
- [Conditional Branch Instructions](#)
- [Miscellaneous Instructions](#)

5.9.1 Data Transfer Instructions

The FPU has two separate register sets: coprocessor general registers and coprocessor control registers. The FPU has a load/store architecture; all computations are done on data held in coprocessor general registers. The control registers are used to control FPU operation. Data is transferred between registers and the rest of the system with dedicated load, store, and move instructions. The transferred data is treated as unformatted binary data; no format conversions are performed, and therefore no IEEE floating point exceptions can occur.

The supported transfer operations are listed in [Table 5.14](#).

Table 5.14 FPU Data Transfer Instructions

Transfer Direction			Data Transferred
FPU general reg	↔	Memory	Word/doubleword load/store
FPU general reg	↔	CPU general reg	Word/doubleword move
FPU control reg	↔	CPU general reg	Word move

5.9.1.1 Data Alignment in Loads, Stores, and Moves

All coprocessor loads and stores operate on naturally-aligned data items. An attempt to load or store to an address that is not naturally aligned for the data item causes an Address Error exception. Regardless of byte-ordering (the endianness), the address of a word or doubleword is the smallest byte address in the object. For a big-endian machine, this is the most-significant byte; for a little-endian machine, this is the least-significant byte (endianness is described in “[Byte Ordering and Endianness](#)” on page 37).

5.9.1.2 Addressing Used in Data Transfer Instructions

The FPU has loads and stores using the same *register+offset* addressing as that used by the CPU. Moreover, for the FPU only, there are load and store instructions using *register+register* addressing.

Tables [5.15](#) through [5.17](#) list the FPU data transfer instructions.

Table 5.15 FPU Loads and Stores Using Register+Offset Address Mode

Mnemonic	Instruction	Defined in MIPS ISA
LDC1	Load Doubleword to Floating Point	microMIPS32
LWC1	Load Word to Floating Point	microMIPS32
SDC1	Store Doubleword to Floating Point	microMIPS32
SWC1	Store Word to Floating Point	microMIPS32

Table 5.16 FPU Loads and Using Register+Register Address Mode

Mnemonic	Instruction	Defined in MIPS ISA
LDXC1	Load Doubleword Indexed to Floating Point	microMIPS32
LUXC1	Load Doubleword Indexed Unaligned to Floating Point	microMIPS32
LWXC1	Load Word Indexed to Floating Point	microMIPS32
SDXC1	Store Doubleword Indexed to Floating Point	microMIPS32
SUXC1	Store Doubleword Indexed Unaligned to Floating Point	microMIPS32
SWXC1	Store Word Indexed to Floating Point	microMIPS32

Table 5.17 FPU Move To and From Instructions

Mnemonic	Instruction	Defined in MIPS ISA
CFC1	Move Control Word From Floating Point	microMIPS32
CTC1	Move Control Word To Floating Point	microMIPS32

Table 5.17 FPU Move To and From Instructions (Continued)

Mnemonic	Instruction	Defined in MIPS ISA
DMFC1	Doubleword Move From Floating Point	microMIPS64
DMTC1	Doubleword Move To Floating Point	microMIPS64
MFC1	Move Word From Floating Point	microMIPS32
MFHC1	Move Word from High Half of Floating Point Register	microMIPS32
MTC1	Move Word To Floating Point	microMIPS32
MTHC1	Move Word to High Half of Floating Point Register	microMIPS32

5.9.2 Arithmetic Instructions

Arithmetic instructions operate on formatted data values. The results of most floating point arithmetic operations meet the IEEE standard specification for accuracy—a result is identical to an infinite-precision result that has been rounded to the specified format, using the current rounding mode. The rounded result differs from the exact result by less than one unit in the least-significant place (ULP).

FPU IEEE-approximate arithmetic operations are listed in [Table 5.18](#).

Table 5.18 FPU IEEE Arithmetic Operations

Mnemonic	Instruction	Defined in MIPS ISA
ABS.fmt	Floating Point Absolute Value (Arithmetic if $FIR_{Has2008}=0$ or $FCSR_{ABS2008}=0$)	microMIPS32
ABS.fmt (PS)	Floating Point Absolute Value (Paired Single) (Arithmetic if $FIR_{Has2008}=0$ or $FCSR_{ABS2008}=0$)	microMIPS32
ADD.fmt	Floating Point Add	microMIPS32
ADD.fmt (PS)	Floating Point Add (Paired Single)	microMIPS32
C.cond.fmt	Floating Point Compare	microMIPS32
C.cond.fmt (PS)	Floating Point Compare (Paired Single)	microMIPS32
DIV.fmt	Floating Point Divide	microMIPS32
MUL.fmt	Floating Point Multiply	microMIPS32
MUL.fmt (PS)	Floating Point Multiply (Paired Single)	microMIPS32
NEG.fmt	Floating Point Negate (Arithmetic if $FIR_{Has2008}=0$ or $FCSR_{ABS2008}=0$)	microMIPS32
NEG.fmt (PS)	Floating Point Negate (Paired Single) (Arithmetic if $FIR_{Has2008}=0$ or $FCSR_{ABS2008}=0$)	microMIPS32
SQRT.fmt	Floating Point Square Root	microMIPS32
SUB.fmt	Floating Point Subtract	microMIPS32
SUB.fmt (PS)	Floating Point Subtract (Paired Single)	microMIPS32

Two operations, Reciprocal Approximation (RECIP) and Reciprocal Square Root Approximation (RSQRT), may be less accurate than the IEEE specification:

- The result of RECIP differs from the exact reciprocal by no more than one ULP.

- The result of RSQRT differs from the exact reciprocal square root by no more than two ULPs.

Within these error limits, the results of these instructions are implementation specific.

A list of FPU-approximate arithmetic operations is given in [Table 5.19](#).

Table 5.19 FPU-Approximate Arithmetic Operations

Mnemonic	Instruction	Defined in MIPS ISA
RECIP.fmt	Floating Point Reciprocal Approximation	microMIPS32
RSQRT.fmt	Floating Point Reciprocal Square Root Approximation	microMIPS32

Four compound-operation instructions perform variations of multiply-accumulate—that is, multiply two operands, accumulate the result to a third operand, and produce a result. These instructions are listed in [Table 5.20](#).

Arithmetic and rounding behavior: The product is rounded according to the current rounding mode prior to the accumulation. The accumulated result is also rounded. This model meets the IEEE-754-1985 accuracy specification; the result is numerically identical to an equivalent computation using a sequence of multiply, add/subtract, or negate instructions. Similarly, exceptions and flags behave as if the operation was implemented with a sequence of multiply, add/subtract and negate instructions. This behavior is often known as “Non-Fused”.

[Table 5.20](#) lists the FPU Multiply-Accumulate arithmetic operations.

Table 5.20 FPU Multiply-Accumulate Arithmetic Operations

Mnemonic	Instruction	Defined in MIPS ISA
MADD.fmt	Floating Point Multiply Add	microMIPS32
MADD.fmt (PS)	Floating Point Multiply Add (Paired Single)	microMIPS32
MSUB.fmt	Floating Point Multiply Subtract	microMIPS32
MSUB.fmt (PS)	Floating Point Multiply Subtract (Paired Single)	microMIPS32
NMADD.fmt	Floating Point Negative Multiply Add	microMIPS32
NMADD.fmt (PS)	Floating Point Negative Multiply Add (Paired Single)	microMIPS32
NMSUB.fmt	Floating Point Negative Multiply Subtract	microMIPS32
NMSUB.fmt (PS)	Floating Point Negative Multiply Subtract (Paired Single)	microMIPS32

5.9.3 Conversion Instructions

These instructions perform conversions between floating point and fixed point data types. Each instruction converts values from a number of operand formats to a particular result format. Some conversion instructions use the rounding mode specified in the *Floating Control/Status* register (*FCSR*), while others specify the rounding mode directly.

[Tables 5.21](#) and [5.22](#) list the FPU conversion instructions according to their rounding mode.

Table 5.21 FPU Conversion Operations Using the *FCSR* Rounding Mode

Mnemonic	Instruction	Defined in MIPS ISA
CVT.D.fmt	Floating Point Convert to Double Floating Point	microMIPS32

Table 5.21 FPU Conversion Operations Using the *FCSR* Rounding Mode

Mnemonic	Instruction	Defined in MIPS ISA
CVT.L.fmt	Floating Point Convert to Long Fixed Point	microMIPS32
CVT.PS.S	Floating Point Convert Pair to Paired Single	microMIPS32
CVT.S.fmt	Floating Point Convert to Single Floating Point	microMIPS32
CVT.S.fmt (<i>PL, PU</i>)	Floating Point Convert to Single Floating Point (Paired Lower, Paired Upper)	microMIPS32
CVT.W.fmt	Floating Point Convert to Word Fixed Point	microMIPS32

Table 5.22 FPU Conversion Operations Using a Directed Rounding Mode

Mnemonic	Instruction	Defined in MIPS ISA
CEIL.L.fmt	Floating Point Ceiling to Long Fixed Point	microMIPS32
CEIL.W.fmt	Floating Point Ceiling to Word Fixed Point	microMIPS32
FLOOR.L.fmt	Floating Point Floor to Long Fixed Point	microMIPS32
FLOOR.W.fmt	Floating Point Floor to Word Fixed Point	microMIPS32
ROUND.L.fmt	Floating Point Round to Long Fixed Point	microMIPS32
ROUND.W.fmt	Floating Point Round to Word Fixed Point	microMIPS32
TRUNC.L.fmt	Floating Point Truncate to Long Fixed Point	microMIPS32
TRUNC.W.fmt	Floating Point Truncate to Word Fixed Point	microMIPS32

5.9.4 Formatted Operand-Value Move Instructions

These instructions all move formatted operand values among FPU general registers. A particular operand type must be moved by the instruction that handles that type. There are four kinds of move instructions:

- Unconditional move
- Instructions which modify the sign bit (ABS.fmt and NEG.fmt when $FCSR_{ABS2008}=1$)
- Conditional move that tests an FPU true/false condition code
- Conditional move that tests a CPU general-purpose register against zero

Conditional move instructions operate in a way that may be unexpected. They always force the value in the destination register to become a value of the format specified in the instruction. If the destination register does not contain an operand of the specified format before the conditional move is executed, the contents become **UNPREDICTABLE**. (For more information, see the individual descriptions of the conditional move instructions in Volume II.)

These instructions are listed in Tables 5.23 through 5.25.

Table 5.23 FPU Formatted Operand Move Instructions

Mnemonic	Instruction	Defined in MIPS ISA
ABS.fmt	Floating Point Absolute Value (Non-Arithmetic if $FCSR_{ABS2008}=1$)	microMIPS32

Table 5.23 FPU Formatted Operand Move Instructions

Mnemonic	Instruction	Defined in MIPS ISA
ABS.fmt (<i>PS</i>)	Floating Point Absolute Value (Paired Single) (Non-Arithmetic if $FCSR_{ABS2008}=1$)	microMIPS32
MOV.fmt	Floating Point Move	microMIPS32
MOV.fmt (<i>PS</i>)	Floating Point Move (Paired Single)	microMIPS32
NEG.fmt	Floating Point Negate (Non-Arithmetic if $FCSR_{ABS2008}=1$)	microMIPS32
NEG.fmt (<i>PS</i>)	Floating Point Negate (Paired Single) (Non-Arithmetic if $FCSR_{ABS2008}=1$)	microMIPS32

Table 5.24 FPU Conditional Move on True/False Instructions

Mnemonic	Instruction	Defined in MIPS ISA
MOVE.fmt	Floating Point Move Conditional on FP False	microMIPS32
MOVE.fmt (<i>PS</i>)	Floating Point Move Conditional on FP False (Paired Single)	microMIPS32
MOVT.fmt	Floating Point Move Conditional on FP True	microMIPS32
MOVT.fmt (<i>PS</i>)	Floating Point Move Conditional on FP True (Paired Single)	microMIPS32

Table 5.25 FPU Conditional Move on Zero/Nonzero Instructions

Mnemonic	Instruction	Defined in MIPS ISA
MOVN.fmt	Floating Point Move Conditional on Nonzero	microMIPS32
MOVN.fmt (<i>PS</i>)	Floating Point Move Conditional on Nonzero (Paired Single)	microMIPS32
MOVZ.fmt	Floating Point Move Conditional on Zero	microMIPS32
MOVZ.fmt (<i>PS</i>)	Floating Point Move Conditional on Zero (Paired Single)	microMIPS32

5.9.5 Conditional Branch Instructions

The FPU has PC-relative conditional branch instructions that test condition codes set by FPU compare instructions (C.cond.fmt).

All branches have an architectural delay of one instruction. When a branch is taken, the instruction immediately following the branch instruction is said to be in the **branch delay slot**, and it is executed before the branch to the target instruction takes place. Conditional branches come in two versions, depending upon how they handle an instruction in the delay slot when the branch is not taken and execution falls through:

- **Branch** instructions execute the instruction in the delay slot.

Branch likely instructions do not execute the instruction in the delay slot if the branch is not taken (they are said to *nullify* the instruction in the delay slot). The microMIPS architectures do not include any branch likely instructions. They are listed here only for historical background.

The MIPS64 Architecture defines eight condition codes for use in compare and branch instructions. For backward compatibility with previous revision of the ISA, condition code bit 0 and condition code bits 1 thru 7 are in discontinuous fields in *FCSR*.

Table 5.26 lists the conditional branch FPU instructions;

Table 5.26 FPU Conditional Branch Instructions

Mnemonic	Instruction	Defined in MIPS ISA
BC1F	Branch on FP False	microMIPS32
BC1T	Branch on FP True	microMIPS32

5.9.6 Miscellaneous Instructions

The MIPS ISA defines various miscellaneous instructions that conditionally move one CPU general register to another, based on an FPU condition code. It also defines an instruction to align a misaligned pair of paired-single values (ALNV.PS) and a quartet of instructions that merge a pair of paired-single values (PLL.PS, PLU.PS, PUL.PS, PUU.PS).

Table 5.27 lists these conditional move instructions.

Table 5.27 CPU Conditional Move on FPU True/False Instructions

Mnemonic	Instruction	Defined in MIPS ISA
ALNV.PS	FP Align Variable	microMIPS32
MOVN.fmt	Move Conditional on FP False	microMIPS32
MOVZ.fmt	Move Conditional on FP True	microMIPS32
PLL.PS	Pair Lower Lower	microMIPS32
PLU.PS	Pair Lower Upper	microMIPS32
PUL.PS	Pair Upper Lower	microMIPS32
PUU.PS	Pair Upper Upper	microMIPS32

5.10 Valid Operands for FPU Instructions

The floating point unit arithmetic, conversion, and operand move instructions operate on formatted values with different precision and range limits and produce formatted values for results. Each representable value in each format has a binary encoding that is read from or stored to memory. The binary encodings of the operand formats are described in Volume II-B: The microMIPS Instruction Set. A conversion instruction specifies the result type in the *function* field; the result of other operations is given in the same format as the operands.

Table 5.28 FPU Operand Formats

Instruction Mnemonic	Size		Data Type
	Name	Bits	
S	single	32	Floating point
D	double	64	Floating point
W	word	32	Fixed point

Table 5.28 FPU Operand Formats

Instruction Mnemonic	Size		Data Type
	Name	Bits	
L	long	64	Fixed point
PS	paired single	64 (2x32)	Floating point

The result of an instruction using operand formats marked **U** in Table 5.31 is not currently specified by this architecture and causes a Reserved Instruction exception.

Table 5.29 Valid Formats for FPU Operations

Mnemonic	Operation	Operand Fmt				
		Float			Fixed	
		S	D	PS	W	L
ABS	Absolute value	•	•	•		
ADD	Add	•	•	•		
C.cond	Floating Point compare	•	•	•		
CEIL.L, (CEIL.W)	Convert to longword (word) fixed point, round toward $+\infty$	•	•			
CVT.D	Convert to double floating point	•			•	•
CVT.L	Convert to longword fixed point	•	•			
CVT.S	Convert to single floating point		•		•	•
CVT. PU, PL	Convert to single floating point (paired upper, paired lower)			•		
CVT.W	Convert to 32-bit fixed point	•	•			
DIV	Divide	•	•			
FLOOR.L, (FLOOR.W)	Convert to longword (word) fixed point, round toward $-\infty$	•	•			
MADD.S	Multiply-Add Single	•				
MADD.D	Multiply-Add Double		•			
MADD.PS	Multiply-Add Paired-Single			•		
MOV	Move Register	•	•	•		
MOVC	FP Move conditional on condition	•	•	•		
MOVN	FP Move conditional on GPR \neq zero	•	•	•		
MOVZ	FP Move conditional on GPR=zero	•	•	•		
MSUB.S	Multiply-Subtract Single	•				
MSUB.D	Multiply-Subtract Double		•			
MSUB.PS	Multiply-Subtract Paired-Single			•		
MUL	Multiply	•	•	•		
NEG	Negate	•	•	•		
NMADD.S	Negative Multiply-Add Single	•				
NMADD.D	Negative Multiply-Add Double		•			

Table 5.29 Valid Formats for FPU Operations (Continued)

Mnemonic	Operation	Operand Fmt				
		Float			Fixed	
		S	D	PS	W	L
NMADD.PS	Negative Multiply-Add Paired-Single			•		
NMSUB.S	Negative Multiply-Subtract Single	•				
NMSUB.D	Negative Multiply-Subtract Double		•			
NMSUB.PS	Negative Multiply-Subtract Paired-Single			•		
PLL, PLU, PUL, PUU	Pair (Lower Lower, Lower Upper, Upper Lower, Upper Upper)			•		
RECIP	Reciprocal Approximation	•	•			
ROUND.L, (ROUND.W)	Convert to longword (word) fixed point, round to nearest/even	•	•			
RSQRT	Reciprocal square root approximation	•	•			
SQRT	Square Root	•	•			
SUB	Subtract	•	•	•		
TRUNC.L, (TRUNC.W)	Convert to longword (word) fixed point, round toward zero	•	•			
Key: • – Valid. Blank - not represented in format encoding						

5.11 FPU Instruction Formats

An FPU instruction is a single 32-bit word. Instruction Formats are shown in Volume II-B: The microMIPS instruction set.

Revision History

In the left hand page margins of this document you may find vertical change bars to note the location of significant changes to this document since its last release. Significant changes are defined as those which you should take note of as you use the MIPS IP. Changes to correct grammar, spelling errors or similar may or may not be noted with change bars. Change bars will be removed for changes which are more than one revision old.

Please note: Limitations on the authoring tools make it difficult to place change bars on changes to figures. Change bars on figure titles are used to denote a potential change in the figure itself.

Revision	Date	Description
0.95	March 12, 2001	External review copy of reorganized and updated architecture documentation.
1.00	August 29, 2002	Update based on all feedback received: <ul style="list-style-type: none"> • Fix bit numbering in FEXR diagram • Clarify the description of the width of FPRs in 32-bit implementations • Correct tag on FIR diagram. • Update the compatibility and subsetting rules to capture the current requirements. • Remove the requirement that a licensee must consult with MIPS Technologies when assigning SPECIAL2 function fields.
1.90	September 1, 2002	Update the specification with the changes due to Release 2 of the Architecture. Changes included in this revision are: <ul style="list-style-type: none"> • The Coprocessor 1 FIR register was updated with new fields and interpretations. • Update architecture and ASE summaries with the new instructions and information introduced by Release 2 of the Architecture.
2.00	June 8, 2003	Continue the update of the specification for Release 2 of the Architecture. Changes included in this revision are: <ul style="list-style-type: none"> • Correct the revision history year for Revision 1.00 (above). It should be 2002, not 2001. • Remove NOR, OR, and XOR from the 2-operand ALU instruction table.
2.50	July 1, 2005	Changes in this revision: <ul style="list-style-type: none"> • Correct the wording of the hidden modes section (see Section 2.2, "Compliance and Subsetting"). • Update all files to FrameMaker 7.1. • Allow shadow sets to be implemented without vectored interrupts or support for an external interrupt controller. In such an implementation, they are software-managed.
2.60	June 25, 2008	<ul style="list-style-type: none"> • COP3 no longer extendable by customer. • Section on Instruction fetches added - 1. fetches & endian-ness 2. fetches & CCA 3. self-modified code
2.61	December 5, 2009	<ul style="list-style-type: none"> • Fixed paragraph numbering between chapters. • FPU chapter didn't make it clear that MADD/MSUB were non-fused.

Revision History

Revision	Date	Description
3.00	March 25, 2010	<ul style="list-style-type: none"> • Changes for microMIPS. • List changes in Release 2.5+ and non-microMIPS changes in Release 3. • List PRA implementation options.
3.01	December 10, 2010	<ul style="list-style-type: none"> • Change Security Classification for microMIPS AFP versions.
3.02	March 06, 2011	<ul style="list-style-type: none"> • There is no persistent interpretation of FPR values between instructions. The interpretation comes from the instruction being executed. • Clarification that the PS format availability is solely defined by the FIR.PS bit.
3.50	September 20, 2012	<ul style="list-style-type: none"> • Mention EVA load, store instructions • Define Architecture version of UCA. • IEEE2008, MAC2008, ABS2008, NAN2008 status bits for FPU. • Mention SegCtl, TLBInv*, EVA in Intro.
5.00	December 14, 2012	<ul style="list-style-type: none"> • R5 changes - mention MSA and VZ modules • R5 change - DSP and MT are now modules • Generated QNAN values - changed to use more common bit patterns
5.01	December 15, 2012	<ul style="list-style-type: none"> • No technical content change: • Updated cover for logos • Updated copyright text.
5.02	April 12, 2013	<ul style="list-style-type: none"> • R5 changes: FR=1 64-bit FPU register model required is required, if floating point is supported. Section 2.1.2.4 MIPSr5 Architecture. Section 2.2 Compliance and Subsetting. Section 2.8.5 FPU Registers. Chapter 5 Overview of the FPU Instruction Set: Section 5.1 Binary Compatibility. Section 5.5 Floating Point register Types. Section 5.5.1 FPU Register Models. • R5 change: if any R5 feature, other features must be R5. E.g. if VZ or MSA is implemented, then if floating point is implemented then FR=1 must be implemented. Section 2.2 Compliance and Subsetting. • R5 change retroactive to R3: removed FCSR.MCA2008 bit: no architectural support for fused multiply add with no intermediate rounding. Section 2.1.2.3 MIPSr3 Architecture. Table 5.4 FIR Register Field Descriptions, HAS2008 bit. Figure 5-12 FCSR register Format: MAC2008 bit removed. Section 5.9.2 Arithmetic Instructions: paragraph titled "Arithmetic and rounding behavior". • R5 change: UFR (User mode FR changing): UFR, UNFR, FIR.UFRP, CTC1 and CFC1 changes. Section 5.6 Floating Point Control Registers (FCRs) - UFR and UNFR FCR numbers; Figure 5-11 FIR Register Format, Table 5.6 FIR Register Field Descriptions - UFRP bit; Section 5.6.2 UFR Register and Section 5.6.3 UNFR Register.

Revision	Date	Description
5.03	August 21, 2013	<ul style="list-style-type: none"> Resolved inconsistencies with regards to the availability of instructions in MIPS32r2: MADD.fmt family (MADD.S, MADD.D, NMADD.S, NMADD.D, MSUB.S, MSUB.D, NMSUB.S, NMSUB.D), RECIP.fmt family (RECIP.S, RECIP.D, RSQRT.S, RSQRT.D), and indexed FP loads and stores (LWXC1, LDXC1, SWXC1, SDXC1). The appendix section A.2 “Instruction Bit Encoding Tables”, shared between Volume I and Volume II of the ARM, was updated, in particular the new upright delta Δ mark is added to Table A.2 “Symbols Used in the Instruction Encoding Tables”, replacing the inverse delta marking ∇ for these instructions. Similar updates made to microMIPS’s corresponding sections. Instruction set descriptions and pseudocode in Volume II, Basic Instruction Set Architecture, updated. These instructions are required in MIPS32r2 if an FPU is implemented. . Misaligned memory access support for MSA: see Volume II, Appendix B “Misaligned Memory Accesses”. Has2008 is is required as of release 5 - Table 5.4, “FIR Register Descriptions”. ABS2008 and NAN2008 fields of Table 5.7 “FCSR RegisterField Descriptions” were optional in release 3 and could be R/W , but as of release 5 are required, read-only, and preset by hardware. FPU FCSR.FS Flush Subnormals / Flush to Zero behaviour is made consistent with MSA behaviour, in MSACSR.FS: Table 5.7, “FCSR Register Field Descriptions”, updated. New section 5.8.1.4 “Alternate Flush to Zero Underflow Handling”. Volume I, Section 2.2 “Compliance ad Subsetting” noted that the L format is required in MIPS FPUs, to be consistent with Table 5.4 “FIR Register Field Definitions” . Noted that UFR and UNFR can only be written with the value 0 from GPR[0]. See section 5.6.5 “User accessible FPU Register model control (UFR, CP1 Control Register 1)” and section 5.6.5 “User accessible Negated FPU Register model control (UNFR, CP1 Control Register 4)”